

Fast Track to THE WEB

The Essentials

Get Connected



How The Web Works



Web Browsers



IM And VoIP



Security



E-mail Clients



Newgroups and RSS



<http://www.opera.com>

Fast Track to The Web

The Essentials

By Team Digit

Credits

The People Behind This Book

EDITORIAL

Robert Sovereign-Smith Assistant Editor
Ram Mohan Rao Writer, Copy Editor
Nimish Chandiramani Copy Editor
Bhaskar Sarma Writer
Asfaq Tapia Writer
R Srikanth Writer
Madhusudan Mukerjee Writer

DESIGN AND LAYOUT

Vijay Padaya, Sivalal S Layout Designers
Sivalal S Cover Design
Harsho Mohan Chattoraj Illustrator

© Jasubhai Digital Media a part of 9.9 Media

Published by Jasubhai Digital Media, a part of 9.9 Media. No part of this book may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior written permission of the publisher.

February 2008

Free with Digit. Not to be sold separately. If you have paid separately for this book, please e-mail the editor at editor@thinkdigit.com along with details of location of purchase, for appropriate action.

Newly-wired?

Some of you might think the time for this little book came and went a while ago; on the other hand, we know there are many of you who will appreciate the following eight chapters, which cover most of what you need to know if you're new to the Internet.

Section I talks about the Internet in general, and Section II about software you use when you're connected. In Section I, then, we talk about how to choose an Internet Service Provider, what plan to choose, what settings to change on your computer, and so on. The next chapter talks about how the Internet works, for those of you who are wondering.

Following that are descriptions and mini-reviews of the most common (and essential) software you'll be using when online, most notably the Internet browser.

Communication software spans three categories, though we haven't put them in a single chapter: Instant Messengers, VoIP software (using which you can talk to people using a computer instead of a telephone), and e-mail clients. RSS readers are in a category of their own; for those who don't know, they're the current way of keeping in touch with the news as it appears on the Internet.

Rounding it off is a chapter on security. For some of you, it might seem like we've been talking about it too much; we should say that one can never talk too much about security, especially to people just recently online (or getting online right now)!

It's a little difficult for most people when they begin to use the Internet for the first time—everything can seem confusing, and not everything makes sense. Things ease up very soon, however, and one gets the hang of it rather quickly. This book is to ease those initial pains—if you have a good idea of the basics, and of the software you're using, that's half the job done!

CONTENTS

Section I	Getting Started	7
Chapter 1	Get Connected	8
1.1	Internet Service Providers	9
1.2	Configuring Windows	16
Chapter 2	How The Web Works	21
2.1	The Structure of the Internet	22
2.2	How Web Site Statistics Are Calculated	32
Section II	Software	42
Chapter 3	Web Browsers	43
3.1	Internet Explorer	44
3.2	Maxthon	50
3.3	Firefox	54
3.4	SeaMonkey	64
3.5	Flock	66
3.6	Opera	70
Chapter 4	IM and VoIP	75
4.1	Chat Clients	76
4.2	Multi-protocol IM Clients	81
4.3	VoIP	87
Chapter 5	Security	92
5.1	Introduction	93
5.2	Terminology	93
5.3	Steps To Complete Security	96
5.4	Anti-virus Software	97
5.5	Firewalls	100
5.6	Anti-spyware	105
5.7	Anti-rootkits	109
5.8	History Cleaners	111
5.9	Anti-phishing Tools	116
5.10	Miscellaneous Tools	117
5.10	Miscellaneous Tools	121
Chapter 6	E-mail Clients	122
6.1	The Various Ways Of Using E-mail	123
6.2	IMAP Vs. POP3	124
6.3	The Choices	125
6.4	Outlook Express (OE)	128
6.5	Windows Live Mail (XP and Vista)	131
6.6	Mozilla Thunderbird version 2	134
6.7	The Bat!	137
6.8	Incredimail Xe	140
Chapter 7	Newsgroups and RSS	142
7.1	Newsgroups	143
7.2	RSS Feed Readers	156

Section I

Getting Started

Get Connected



You have a computer at your disposal and are keen to use it to its full potential. You want to improve your research, contact people from all over the world, make more impressive presentations, update your knowledge, use the latest data, express your views, or listen to music. What you need to do is to connect your computer to the Internet.

The Internet is nothing but a network of networks, a Web of computers linked to each other all over the whole wide world (that is not what “www” stands for!) This chapter will tell you just what you’ll have to do to access the Web.

Checking out Web sites is sometimes called “surfing.” Before you start surfing, you will, naturally, need a surfboard—in this case, an ISP (Internet Service Provider).

1.1 Internet Service Providers

1.1.1 What Is An ISP?

ISPs are sometimes also referred to as Internet Access Providers (IAP). An Internet Service Provider is an organisation that provides access to the Internet and related services.

An ISP might provide dial-up access, a cable connection, an ADSL connection, or other type of Internet access. In India, there are more than 150 ISPs (including Icenet, BSNL, Airtel, 24Online, and You Telecom).

1.1.2 Kinds Of Connections

ISPs provide different kinds of connections which allow you to access the Internet, each of which has advantages and drawbacks.

a) Dialup accounts:

This is the slowest kind of connection currently available. Also, your telephone line stays in use until you terminate the connection, implying that if your telephone bill is calculated according to the time of use, you’re going to have one hefty bill coming at you at the end of the month.

Dialup means you won’t really be able to do audio/video downloads unless you have the patience of Sisyphus. Even images can take a lot of time to be downloaded (remember that Web pages are getting prettier and more image-heavy).

Today, a dialup service is the resort of a desperate customer, who has by and large no other means of connecting to the Web.

b) **Broadband:** includes several transmission technologies, such as:

- Digital Subscriber Line (DSL)
- Cable
- Wireless
- Satellite

1. **Digital Subscriber Line (DSL)** is a technology that provides a dedicated digital circuit between a residence and a telephone company's central office, allowing high-speed data transport over existing telephone lines. Typically, there is a modem on either end; one modem is located at the central office, and the other at the customer's site. There are several types of DSL including ADSL, HDSL, IDSL, SDSL, and VDSL, but of these, the most common, at least in the UK and India is ADSL.

ADSL (Asymmetric Digital Subscriber Line) is designed to deliver more bandwidth downstream (from the central office to the customer site) than upstream (that's why the word "asymmetric").

SDSL, common in Europe, is the *Symmetric Digital Subscriber Line*, a technology that allows the same data rates for upstream as well as downstream traffic, over existing copper telephone lines. However, it doesn't allow simultaneously voice connections on the same wires.

2. Cable:

As millions of homes already have cable TV connections, cable TV service providers simply equip the customer with a cable modem and provide access to the Internet over cable TV lines.

The coaxial cable allows for much greater bandwidth than telephone lines, so a cable modem is considerably faster than DSL (about 3 Mbps).

3. Wireless Internet Access

Wireless communication is nothing but the transferring of data over a particular distance without using the long, irritating, tubular things called wires.

Certain ISPs offer access to the Internet without your first having to plug in a wire to a jack. Such services may use different types of wireless technologies:

- **GSM:** Global System for Mobile communication, which compresses information and then transfers it over a single channel.
- **CDMA:** Code Division Multiple Access transfers data by distributing it among different available radio channels.
- **Wi-Fi:** A common wireless network, broadcasting radio waves that can be caught by Wi-Fi receivers in individual computers. The Wi-Fi service operates by setting up Wireless “access points” that connect a group of wireless devices.
- **WAP:** (Wireless Access Protocol) is a standard that allows cellular / wireless devices to access the Internet. The main application of this standard is to allow cellular phones or PDAs (Personal Digital Assistants) or Notebook computers to get access to the Internet, again, without having to plug a wire in.

4. **Satellite Internet Service** is ideal for those who live in rural areas where high speed is not yet available. Quite akin to the way in which we use DTH to receive television content today, it needs individual satellite dishes to be set up which send and receive signals via satellite.

1.1.3 Broadband Or High-speed Internet?

We agree it's confusing for the layman. Should you go for a broadband connection or should you opt for a high-speed connection?

In reality, both terms refer to the same thing: a network connection that can accommodate a large number of frequencies, dividing them into channels, and transferring data at speeds much faster than dialup services do. This sort of network connection allows more data to be transmitted simultaneously through a single connection—broader bandwidth—and this is precisely why the service is called “broadband.”

However, sometimes ISPs advertise “high-speed dialup Internet” or “Accelerated dialup” services, which is not broadband at all, but simply your traditional dialup facility that uses a shortened logon process (known as the handshake) and compresses data differently to speed things up. Besides, it also uses what is known as an “Acceleration Server,” which uses a broadband connection to search the Internet for the server that hosts the site you’re looking for and sends it to your computer over your telephone line. The

transfer rate of the network, measured as the number of bits per second (known as the “throughput”) is the same, and, as mentioned earlier, the telephone line would remain occupied as long as you continue to access the Net.

1.1.4 Choosing An Internet Service Provider

Perhaps the most dicey part of accessing the Internet is finding the right ISP. The first decision you need to make is whether to go with a dialup ISP or a broadband ISP such as a DSL provider or satellite or cable Internet service. We thought it would help if you could have an overview of the differences between dialup and broadband services, and so, here's a table:

Whichever you opt for, when you're choosing an ISP there are few fundamentals you should consider:

- 1. Local phone number:** ISPs typically need you to dial a number that your computer must contact to connect to the Internet. If this number is not local but a long-distance one, you could spend a lot more than you intend to.
- 2. Cost:** Costs that you will incur in using an ISP's services vary from one provider to another, and is usually the factor that will most influence your choice. Each ISP gives a certain number of hours per month which you can use. Charges differ according to the kinds of service you need and the quantum of data you give or receive. Off-peak hours are sometimes free, while extra charges are applied for every extra hour (or MB) you use. Some ISPs charge a one-time fee for setup and then a monthly rental. We advise you to first determine the kind of usage you have in mind. If you are an occasional user and don't think you'll need the Net much, just settle for good old dialup. But if you need to send or receive a lot of data, you could choose from the variety of schemes that are offered by ISPs that let you connect through cable or DSL.
- 3. Mailboxes:** ISPs often allow you to have multiple e-mail addresses for your Internet account. This is sometimes free and can be availed of.

4. **Software:** ISPs usually provide a CD-ROM with software that you must use to connect to and use the Internet. This disc already has the software you need.
5. **Technical support:** It would be ideal if you could opt for an ISP that gives you a helpline number that you could call 24 x 7 if you encounter any glitches (and believe us, there could be plenty). Often, these help-lines take ages to actually reach an operator on the other end and however unimportant that seems, think it out—or better, call the number and try it out even before you actually subscribe to the service.
6. **Bandwidth and speed:** How many channels does your ISP claim to manage simultaneously over a single line? How much data do they claim to be able to transmit per second and how much, in reality does get transported?
7. **Accessibility:** if you're using a modem, the busy signal is possibly going to be one of your daily hindrances in attempting to access the Internet. You could waste a lot of time dialling and redialling. The best way to check this out would be to ask the Internet users in the same area whether they face difficulties in getting connected to the ISP you are intending to subscribe to.
8. **Latency:** It would be to your benefit if you could get an idea of the latency (delay in data transmission) problems that occur if you use a particular ISP. Latency is basically the amount of time a packet of data needs to reach from one point in the network to another.

1.1.5 Identifying ISPs In Your Area Or City:

1. Ask friends. Actually, even ask your enemies. There's nothing like getting the low-down straight from the horse's mouth. People who already use the Internet would be ideal to tell you about their experiences with particular ISPs and help you to work out a comparative evaluation of the most optimal one in your area.

Dial Up		Broadband	
Pluses	Minuses	Pluses	Minuses
Accessible from anywhere you've got a telephone line	Slow	Really fast	Sometimes erratic if it's a DSL on account of traffic, (i.e. a large number of users) at a particular time.
Needs only a traditional 56 kbps modem that most computers have	Limited bandwidth. Very cumbersome for gaming, audio, and video files	Plenty of bandwidth. Go gaming.	ADSL services sometimes make uploading bulky files a pain in the neck
Keeps the telephone occupied and so keeps away unwanted calls (this is only for those of you who really hate socialising on the telephone)	Keeps the telephone occupied preventing you from accessing the Web and making that important phone call at the same time	Keeps your telephone line free for use. Excellent for business, long-distance assignments, etc. where you need your Net connection to be operating while discussing the job with your associates on the telephone.	DSL lines which use your copper wire telephone lines to access the Net often interfere with your telephone conversation with a squeaky background noise if the Internet is accessed over the same line at the same time.
Cheaper. Doesn't demand any extra charges like monthly rentals, etc.	Considering that downloading or uploading a file could take a much longer time, the call charges shoot up, sometimes making dialup access a very expensive proposition. If you need to use the Internet for several hours a day, the phone bills will skyrocket.	Pay a fixed monthly/yearly rate and get a limited amount of space that you can use per month (usually 1 GB), or alternatively, opt for schemes that offer unlimited usage.	If you exceed the allotted limit, the charges increase per MB of access or download you do.
You can log in and log out as per your need	The process of logging in—or accessing the ISP server—sometimes takes a very, very long time	Keep it on 24x7	You might end up sitting at the computer too long, browsing unnecessary sites, winning too many international gaming contests, getting a crick in the neck, carpal tunnel syndrome, etc.

2. Check out their Web sites. These same people who already use the Internet, or the cyber-cafes in your vicinity would be more than willing to give you a few minutes to access the Web sites of popular ISPs. Most ISPs have their tariffs and schemes detailed on their sites. Here are some of the URLs that would be of help:

- <http://airtel-broadband.com>
- <http://broadband.sify.com>
- <http://mumbai.mtnl.Net.in/Internetservices/index.html>
- www.bsnl.co.in/service/dataone.htm
- www.tataindicombroadband.com
- www.rcom.co.in/Webapp/Communications/rcom/Netconnect/rm_netconnect.jsp

3. Look for the ads of the ISPs and read them carefully. If possible, visit their local offices and get hold of a brochure of their services. Keep the list of questions regarding basics of choosing ISPs in §1.1.4 above in mind and ask where necessary.

Each ISP usually offers separate time-based (a certain number of hours), volume-based (a limited available MB of space), or unlimited plans according to volume. To make your task a trifle easier, here's a comparative list of details of the unlimited plans offered by some popular ISPs in India. As there are a wide variety of options available with every ISP today, we have limited this list to the least expensive unlimited plans on offer today in Mumbai city.

Note that these details have been chosen to compare the cheapest tariffs of broadband services available in the Mumbai region. This is by no means a cross-section of all ISPs in the country. Each of the above ISPs has several other attractive schemes geared towards fulfilling the needs of businesses as well as individual consumers, the listing of which would go beyond the scope of this *Fast Track*. Several also offer combo services, which combine phone line usage along with broadband access to the Internet.

Name of the ISP	Scheme	Speed	Usage permitted	Technology used	Tariff (Rs per month)
Sify Broadband Hi-Speed Internet	Hi-speed Plan	128 kbps	Unlimited	Fixed wireless	711
Airtel Broadband Services	Unlimited 499	64 kbps	Unlimited	DSL	499
Reliance Broadnet	Flat Fee Plan	75 kbps	Unlimited	ADSL	400
Tata Indicom Broadband	Infinity 128	128 kbps	Unlimited	DSL	700
BSNL DataOne Icenet 24Online-	Home UL 900 Plus	256 kbps	Unlimited	ADSL	900
Internet on Cable	UNL 1	64 kbps	Unlimited	Internet on cable	600
You Telecom	UN450 UN650	96kbps 192 kbps	Unlimited	Fibre-optic network	450/- p.m. + 60/- as rentals 650 p.m. + 60/- as rentals

*Conditions apply

1.2 Configuring Windows

Right, so now you've finally gone and got yourself registered with an ISP and you're all click-happy to hit the Net. First, however, if your ISP doesn't come and do it for you, you'll have to configure Windows to accept this new addition to your system. In short, you'll have to go (con)figure.

1.2.1 Figuring Out Configuring

When installing a new device or program on your computer, you may sometimes need to configure it. In this case, after your ISP connects the requisite hardware (modem, cable, etc.), you'll have to set the basic parameters that will allow your system to smoothly interact with the ISP to give you a hassle-free access to Net. These days, thanks to new technologies, it's usually just plug-and-play. Yet, for all you guys who don't want to bother the ISP service men too much, here's how you could do it.

If you're working on a network, prior to setting up any connections at all, you'll need to tell your computer the IP address you've been allotted. If the ISP has given you a permanent one (for a detailed explanation of what an IP address is, take a peek at Chapter 2). Your network administrator would be the right guy to ask for the IP address already allotted to your particular network. If, however, your ISP provides you with a fresh IP address each time you log in—and this is how most private broadband connections prefer it—you'll have to keep the settings to automatic. Here's how:

Step 1: Right click on My Network Places on the Desktop and click on Properties. This opens the Network Connections Window.

Step 2: Right click on the LAN icon and again choose Properties. This takes you to the Local Area Network Properties window, which gives you details regarding which modem you're using, and what items the connection uses.

Step 3: Double-click on the line that says Internet Protocol (TCP/IP). If your ISP provides you with specific IP addresses and DNS addresses, fill them in the exact same way as they gave it to you. This will enable your ISP's computer to identify you as a particular user every time you access the service.

Step 4: With most broadband connections today, simply choose the option that says Obtain an IP address automatically. This is also the default setting in windows.

In case your computer is part of more than one network, you'll have to give alternative configurations too. This can be done by clicking on Alternative Configuration.

The next step is to specify the IP address of the ISP. For this you'll have to use the New Connection Wizard.

1.2.2 Make An Internet Connection In Windows XP

If you want to set up a connection to the Internet—or any other network, for that matter—Microsoft has provided the user-friendly New Connection Wizard. To use this, open Network Connections: Start > Control Panel > Network Connections.

1. Under Network Tasks, click Create a new connection, then click Next.
2. To choose the type connection you want, just click “Connect to the Internet” and click Next.
3. If you already have an account with an ISP, click Set up my connection manually and then click Next.

Set up the connection manually

- If you are connecting to your ISP using a standard 28.8 Kbps, 56 Kbps, or ISDN modem, click Connect using a dial-up modem, click Next, and follow the instructions in the Wizard.
- If your DSL or cable modem ISP connection requires a user-name and password, click Connect using a broadband connection that requires a user name and password, click Next, then follow the instructions in the Wizard.
- If your DSL or cable modem ISP connection is always on and does not require you to type a username and password, click Connect using a broadband connection that is always on, click Next, and then click Finish.



If you have a LAN connection or a cable modem to access the Internet, use the “Always on” option

1.2.3 Share Your Connection

The great thing about Internet Connection Sharing on Windows XP is that you can use networked computers to share a single connection to the Internet. In short, if you already have a group of

connected computers—a Local Area Network, for example—you could take a single Internet connection from an ISP and be able to access the Internet from all the computers in the network.

On The Host Computer

Open the Network Connections from the Control Panel. Right-click the connection that you use to connect to the Internet. For example, if you connect to the Internet using a modem, right-click the connection that you want under dial-up and then go to Properties > Advanced > Settings. Allow other network users to connect through this computer's Internet connection by turning off the firewall.

If you're sharing a dialup connection, select Establish a dial-up connection whenever a computer on my network attempts to access the Internet if you want to permit your computer to automatically connect to the Internet.

When Internet Connection Sharing is enabled, your LAN adapter will be set to use IP address 192.168.0.1. Your computer may lose connectivity with other computers on your network. If these other computers have static IP addresses, it is a good idea to set them to obtain their IP addresses automatically.

The connection to the Internet is shared to other computers on the LAN. The network adapter that is connected to the LAN is configured with a static IP address of 192.168.0.1 and a subnet mask of 255.255.255.0.

On The Client Computer

To connect to the Internet using the shared connection, you must confirm the LAN adapter IP configuration, and then configure the client computer. Follow these steps:

Log on to the client computer as Administrator or as Owner. Go to



You can select which users can connect to your computer

Start > Control Panel > Network Connections. Right-click on Local Area Connection, and then click Properties.

Under the General tab, double-click Internet Protocol (TCP/IP) in the This connection uses the following items list.

In the Internet Protocol (TCP / IP) Properties dialog box, click Obtain an IP address automatically (if it is not already selected), then click OK.

An example of the IP address, subnet mask, and default gateway is given below:

IP Address: 192.168.0.2

Subnet mask: 255.255.255.0

Default gateway: 192.168.0.1

To configure the client computer to use the shared Internet connection, follow these steps:

Left click through Start > Control Panel > Network and Internet Connections > Internet Options. In the Internet Properties dialog box, go to the Connections tab and then the Setup button. Use the good old New Connection Wizard to set up your connection.

When you now start Microsoft Internet Explorer, the client computer will try to connect to the Internet by using the host computer's shared Internet connection.



Click on "Guest" to enable the computer to access information from another computer (the host)

Summing It Up

As is evident from the steps above, the Operating System installed on your computer has made it as easy as possible to perform the initial actions that will allow you to establish the connection you need to link your computer to an ISP. The options in every window are few, and the Next and Finish buttons make this link-up easy enough for kids to do. Once done, you have entered the world of the Internet. The next chapter will tell you how exactly this world works.

How The Web Works



The Internet is not a single network. It is a vast, earth-spanning network of networks. As you're reading this, dozens of networks and PCs are in the process of hooking themselves up to this international network, extending it almost infinitely.

You might just be pining, craving, desperately wishing that somebody would come forward and explain how this whole blessed system works. And that signifies the first step towards geekhood—and if you read on, this chapter could expedite the process.

2.1 The Structure of the Internet

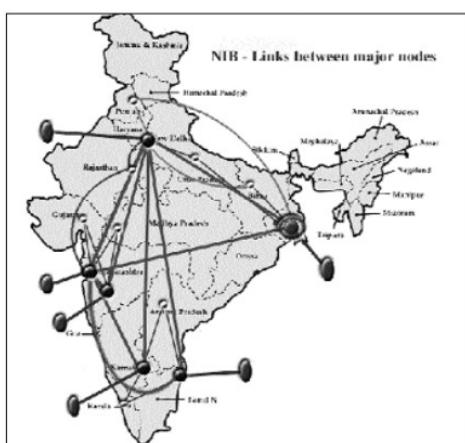
2.1.1 Webs In General

Spiders made webs before we did, you know. At first, a spider makes a bridge line—the first thread that connects two solid points, like two branches or walls or trees. He then reinforces this thread with a few more threads. This primary line is what holds the entire web together. He then proceeds to make a Y-shaped radial—a framework—and then the spokes that constitute a web.

If you thought it was just some bright copywriter's idea to call the Internet a "web," think again. There are some remarkable similarities between how spiders make their webs and how the Internet takes shape. The primary resemblance is what is known as the Internet "backbone."

2.1.2 The Central—But Not Really Nervous—System

Akin to the spider's bridge line is what is called a "backbone," a super-fast network that goes all across the world, usually from one metropolitan area to another. Mega-ISPs run these lines that can transport data at approximately 45 MBps (on T3 lines) and are linked at specified interconnection points (called "national access points" and sometimes known as "regional nerve centres"). All regional networks are connected to each other by high-speed backbones, which are basically connections that can transport data at very high speeds. When data is sent from one regional network to another, it is first sent to the above-mentioned NAP which in turn routes it to a backbone. The ISP that connects your computer to the Internet is a local one that uses routers to send and receive data via this backbone.



The BSNL National Internet Backbone (NIB)

2.1.3 About A Router

A router (the “rou” in “router” rhymes with “cow”) is an electronic device which comes in a little box, and which connects one network with another (for example, a Local Area Network with a Wide Area Network). Routers do most of the task of directing traffic on the Internet by sending the request from your network to another while simultaneously preventing other unauthorised users to access your network (this latter task is done with the help of a firewall, a software program that intercepts and filters unwanted packets of data that are either trying to enter your network or access the resources of your network). A router examines each packet of data that travels across the Internet, figures out where it is headed, and “routes” it towards its destination in the most efficient way possible. Usually, the packet of data goes to another router and another, and another—until it reaches the destination. Each router that a packet of data passes through to reach its destination is considered a “hop.” A router must depend, however, on gateways.

2.1.4 Gateways And Bridges

Just when you—finally—figured out that “windows” aren’t something you open to let fresh air in, these computer guys tell you to start using “gateways,” too. And you got it right this time—a “gateway” is not something you open to let people in.

A gateway is a network point—a node—that acts as an entrance to another network. It is basically a device that allows the exchange of data by translating between the protocols of two different networks or computers on the connected networks. In the case of your residential connection, the gateway is your ISP. Gateways are different from Bridges in the sense that a bridge only links LANs together, sending data to another LAN while keeping local data within its own network.

Along the way, the data we transact—apart from passing through several routers, gateways, and backbones—is also dependent on an indispensable device called a Server. Servers are of different types, and they perform several important functions.

2.1.5 Servers

A server is basically a program that runs on a machine and provides a specific service to all machines that are connected to it. The server has a specific set of programs that allows the connected computers (called clients) to make different kinds of requests to it.

Think of a server as a waiter in a restaurant standing by to give you all he can accommodate on a tray. If, however, he doesn't have the item you wish to consume, he re-routes your order to another waiter (with different goodies) who happens to be waiting at another table. This other waiter passes on the object of desire via other waiters so it finally ends up on your plate.

A Web server is a machine that hosts Web sites and allows Internet users (clients) to access these sites. A Web server is basically the kind of server that holds static companies and just waits around for requests to come around from other machines on the Internet. Once it receives a request, it allows access to the data requested.

Web servers are the principal part of the Web, and every URL you request usually ends up at a Web server—unless a proxy server decided that it can help speed up the process by sending you the data itself.

- A file server is a machine that keeps files and allows clients to upload and download them from it.
- A proxy server sits between a client application, such as a Web browser, and a real server. It intercepts all requests to the real server, sees if it can fulfil the request from any of the recent requests, and if it can't, forwards the request to the server that actually hosts the Web site.
- A groupware server is software designed to enable users to collaborate, regardless of location, via the Internet or a corporate intranet and to work together in a virtual atmosphere.
- Mail servers move and store mail all across the Internet, and you have been “served” by one when you've checked your mail.
- A chat server allows you to interact with other users on the network in real-time.
- Integral to the process of sending data across the Internet is the Domain Name Server (“the DNS”).

2.1.5 The Process Explained—Roughly

To give you a rough sketch of a smooth process, this is what happens when you log on to the Net and type in a URL. The browser

you're using—Internet Explorer, Mozilla Firefox, Safari, Seamonkey, or anything else—is called a client, which sends a “request” to access the URL you typed in (say www.yahoo.co.in) to your ISP. The ISP forwards the request to another (like a file on the table of a typical government office goes to another table) through the vBNS (the Very high-speed Backbone Network Services) and so on, until it reaches the target where the destination is stored—in this case, the Web site you requested. This “host” computer, or “Web host,” sends back the packet of data you requested through much the same channels in much the same way, hopping routers on the way, until it reaches your ISP. The ISP then sends the data down to your computer and the Yahoo! India homepage appears on your screen.

Simple? No. Actually, *mucho, mucho* complicated.

2.1.6 The Process Explained—Toughly

We'll now try to figure out how the Web works by digging a little deeper and see what exactly happens to get your favourite Web page on your screen.

Let's say you typed in the URL of a Web site. The Uniform Resource Locator, for your convenience, is in letters and numbers, in recognisable words, so that you can identify and recall it at will. The computers aren't human... yet. So, out of our deep sympathy for the dumb machines, we humans have conceded to talk to computers in their language, which is nothing but the language of zeroes and ones—“bits.” If all computers used this binary language in random ways, exchange of meaningful data would be impossible. To ensure a shared language between two computing machines, we created “protocols.”

2.1.7 Protocols

A protocol is a set of rules that enables the exchange of information between two computers, regardless of whether they run on different operating systems. These rules strictly state the format and procedure to be followed while transmitting data between two computers or two networks. A protocol basically gives the “grammar” to

their “speech.” Certain key protocols you should know about are:

TCP: Transmission Control Protocol, used for the reliable transmission of data over a network.

IP: The Internet protocol is the set of rules that dictates how exactly packets of data should be transmitted over the Internet. IP standardises the way machines over the Internet (or any IP network) forward or route their packets based on their IP addresses.

HTTP: Hypertext Transfer Protocol, used for transmitting and displaying information in the form of Web pages on browsers. This is the language used when your Web browser talks to the Web server.

FTP: File Transfer Protocol, used for file transfer (uploading and downloading) over the Internet. Both FTP and HTTP are based on TCP/IP. FTP is still often used as a means of downloading large files.

SMTP: Simple Mail Transfer Protocol, used for e-mail.

Ethernet: Used for data transmission over a LAN (Local Area Network).

Wi-Fi: The wireless version of Ethernet.

Of all these protocols, in order to understand how the Internet works, what we really need to know is the IP and the TCP.

2.1.8 TCP/IP

Every device, every computer, that is part of the Internet, is allocated an address, called an Internet address, or Internet Protocol address (IP address). However, it isn't like our home address or anything because, as we mentioned earlier, computers only understand numbers. So the IP address is a series of numbers.

The current protocol is called IPv4. Each IP address is actually just four numbers each ranging from 0 to 255 (each of which is called an “octet”) and separated by decimal points (called dots). IP addresses are, therefore in the format xxx.xxx.xxx.xxx where each xxx could be any number from 0 to 255. So an IP address could look something like 161.184.138.36. This address is essential during every exchange of data on the Internet, because it identifies the client computer (which makes the request, and to which the data must be sent) and the destination computer (to which the request must reach). Nothing, therefore, can happen on the Internet unless your computer is first

assigned an IP address.

Say you've specified the IP address of a Web site you wish to access. A server usually has a static IP address. However, your home PC, if you're connecting to the Internet through a modem, is usually assigned a different IP address by the ISP each time you dial in. This IP address is unique to you as a user, but only for as long as the session lasts. The reason for this system is that ISPs have to deal with a large number of requests at the same time. Besides, at any given time, the number of users of the ISP server is limited. To assign permanent IP addresses to retail users wouldn't really be required; instead, ISPs simply allot any IP address that is not being used at the time to each user who has just dialled in. This system enables the ISP to get along with fewer IP addresses for the number of users who are connected simultaneously.

TCP breaks down and reassembles packets of data. IP ensures that the packets are sent to the right destination. The TCP/IP combination is used because the Internet is a "packet-switched network." In a "circuit-switched network," in contrast, once a connection is made, that specific part of the network is exclusively used only for that connection. In a packet-switched network, the connection between the sender and the receiver is not single and unbroken. When information is sent, it is broken into small packets and sent over many different routes at the same time and then put together in order at the other end once the packets of data reach the destination.

2.1.9 How TCP/IP works

Data sent across the Internet is broken up into packets of less than about 1,500 characters each. Each packet is given a "header," which holds the information needed to put the packets back in right order again. To each header is added a "checksum"—basically a number that checks whether the precise amount of data in the packet has been received or not. Each packet is put into what is called an IP "envelope"—a packet of data with a common address and other specifications. Once the packets go through the routers to their destination, TCP, at the receiver's end, calculates a checksum for each packet and then compares it with the checksum that

was sent in the packet. If these don't match, the computer figures that the data has been corrupted en route, discards the packet, and requests the sender to re-send the packet. TCP then assembles them into their original coherent form to allow the receiving computer to make sense of it.

None of this would be possible, however, without a DNS server acting as mediator.

2.1.10 The Domain Name System

Most of us would find it considerably difficult to remember hundreds of strings of numbers of all the sites we visit. So, for humans, the address is in letters (the URL). The domain name is the root identifier of a Web site on the Internet.

When you want to contact a location on the Internet—say a Web site—you'd type in an address like www.yahoo.com. IP uses (a) the Internet address information and (b) the Domain Name System to understand and deliver your mail or request from one computer to another: the letters www.yahoo.com are translated into an IP address, which is, in this case, 87.248.113.14. In fact, if you typed in <http://87.248.113.14>, your screen would show you the Yahoo! site, and you would make the job a microsecond faster and a trifle easier for the computers. This translation is accomplished by the DNS server.

Before we understand what the DNS server does, let's get a grip of how the DNS works.

2.1.11 How The Domain Name System Works

In order to make efficient and hassle-free the translation from the plain-English www.yahoo.com into numbers that computers can understand, the Internet has been organised into a number of major domains. Major domains are represented by the letters at the end of the English address. These are called top-level domains (TLD) or zones, and they are either two or three letters long. The three-letter zones indicate the type of organisation that owns the domain, that is, whether it's a commercial institution, ISP, non-governmental organisation or academic establishment, etc. All Internet domains are registered with private companies called Internet Registrars, in co-ordination with the InterNIC (Internet

Network Information Center). You probably are most familiar with the .com TLD. Here's a list:

In October 1998, the Internet Corporation for Assigned Names and Numbers (ICANN) was formed, and took over the task of man-

Top-level Domains		
Generic Top Level Domain (gTLD)	Description	Example
.com	For commercial organisations or individuals doing business	www.thinkdigit.com
.edu	For educational domains	www.amity.edu
.net	For networks, Internet service providers, and other network-related companies	www.att.net
.org	For non-governmental organisations and non-commercial or non-profit institutions	www.cry.org
.mil	For the United States military	www.army.mil

aging the domain names and other responsibilities that were fulfilled by the InterNIC. ICANN announced on November 16, 2000, the following seven new gTLDs:

Besides, additional two-letter domain names have been assigned to identify domains belonging to particular countries

gTLD	Description
.aero	Aviation-related businesses
.biz	Businesses (to take the pressure off the.com TLD)
.coop	Co-operatives (for the same reason as above)
.info	For sites that only provide information
.museum	Museums
.name	individuals' real or screen names
.pro	Professionals

(outside the US)—.au, .in and so on.

2.1.12 Understanding Domains

The DNS was developed by Sun Microsystems in the early 1980s,

and has been the addressing system on the Internet since. What the DNS does is to establish a hierarchy of domains, which are groups of computers on the Internet. The DNS gives each computer on the Net an address, or domain name, using easily-recognisable letters and words instead of numbers. Each top-level domain maintains the lists and addresses of the second-level domain, and so on. The letters or word you see before the dot of the gTLD is called the “second-level domain” (like “google” in “www.google.com”). The second-level domain is usually chosen by the organisation that owns the computer.

As an example of how the DNS and domains work, let us consider the example of www.thinkdigit.com. The top domain here is .com, which stands for commercial organisation (not governmental, educational, etc.). The www below the sub-domain name is by far the most commonly-used, as it refers to the organisation’s Web server, that is, the computer that stores all the organisation’s Web pages.

What you pride as your e-mail ID is constructed in a similar way. Take the example xyz@yahoo.co.in. The first part of the address—to the left of the @ sign—is the username. This is what the person called xyz—or who would like to be identified as xyz—will use to log in to his Internet account. The second part of the address—the letters to the right of the @ sign—is the domain name or hostname. This identifies the specific computer where the person has an Internet e-mail account (e.g. a mail server). If the e-mail account is located at the person’s workplace, this host or domain name would be the same as the name of the company (like xyz@bbc.co.uk). Else, it will reflect where the person has an Internet account—an ISP or online service. So, a closer look at the example xyz@yahoo.co.in when read from right to left tells every computer on the Internet to route data to...

The zone in which the host server is located (.in = India)

The nature of the organisation (.co = commercial)

The service that provides the e-mail service (yahoo)

The username of the user (xyz)

2.1.13 The Process explained, but shortly

Let’s move on to analyse the situation a little deeper. How exactly

does the Internet decipher which computer you want to contact to receive data from or send data to? We'll explain all...

Step 1: When a specific URL needs to be contacted, the network needs to match the URL to the actual IP address (the alphabetical gobbledegook to the safe, meaningful language of numbers). Two things could happen now:

- Your Web browser first goes to the local name server (a DNS server) maintained by your ISP or company, and checks if the IP address is on the same network as you. If it finds it right there, it resolves the URL with the corresponding IP address, and sends the true IP address to your computer.
- Or, the information you have requested is not on your local network. The local name server will try to get the information from some other DNS server on the Internet. The local name server contacts the root domain server, which supplies the local name server with details about which Primary Name Server and Secondary Name Server actually hosts the URL you requested. The local name server (the DNS server of your ISP) now contacts the primary name server. If it still can't find the URL there, it contacts the secondary name server. It now sends the information about the true IP address of the URL back to your browser.

Step 2: Either way, your Web browser now has the IP address of the URL you want to access and can use it to contact the proper site.

Step 3: The request is broken up into IP packets by TCP.

Step 4: These packets go to an internal router, a router that may not be on the Internet at all but in your own network (LAN or WAN)—which checks out the address.

Step 5: The internal router now figures out whether the request is targeted to another computer on the same network or one outside the network. If it finds the destination on the same network, it forwards your request to it. Or else it directs your request to an external router on the Internet.

Step 6: The router usually uses “dynamic routing” and allows the packets of data to have multiple routes to reach the destination (the host computer). This enables the router to change the way in which the information can be routed if there are any glitches in the system preventing smooth passage of data, such as traffic or hardware breakdowns.

Step 7: The request reaches the host computer, which runs server software that enables the host to separate the actual request from all the scores of packets of data received.

Step 8: The Web host now sends back to the client PC (from where the request originated) the specific Web page. If the request is a search, then the host computer executes a database search and sends back the results of the search in the form of a Web page. Databases and other such applications on the computer are managed through a common gateway interface (CGI), which forwards the request to a separate database server if necessary.

Step 9: From here on, the process reverses, in a way. The packets of data sent by the Web host are routed back to the client PC through routers and gateways.

Step 10: The client PC (your computer) receives the packets of data, makes sense out of it all using TCP/IP and more, and bingo—there's the Web page on the screen!

This whole process, which we've laboriously waded through, happens in just a few seconds.

2.2 How Web Site Statistics Are Calculated

A lot of nefarious activities occur in the virtual world, and the reason for this is the popular impression that the Web is a place where you can remain anonymous. This, however, is a myth.

While things like your name and residential address can remain secret while you're surfing the Internet (unless you go ahead and volunteer the information), most Web sites collect information about the visitors to their sites while they're sending their Web pages to your PC. This is known as "Web tracking," "Web server analysis," or the reporting of Web site statistics. Web Tracking is, basically, the process of using Web site statistics and analysis tools to determine who is visiting a particular Web site, how they got there, how long they hang around, and what they do once they arrive—in short, the study of information-seeking behaviour. In order to do so, the Web sites usually incorporate tracking software—or Traffic Reporting Software—within the programming content.

2.2.1 Vital Statistics

Now why would anyone want to spend time, money, and other resources on visitors to a Web site? For the same reason that a shopkeeper wants to know the number and the details of visitors to his shop; for the same reason that everyone from little businesses to major companies spend fortunes on surveys and market research. After all, a Web site is a product, and like every manufacturer of a product—or one of its sellers—one should be able to modify, improve, and remodel one's product according to the tastes of the people you hope will buy it.

2.2.2 Advantages of Web tracking

- Tracking the kind of traffic a Web site receives would help the owner of the site to identify areas for database or Web site maintenance. For example, the number of failed searches for a particular term on the Web site could reveal data entry errors that need to be rectified.
- Information so gathered could be used to offer feedback to users on their use of the information retrieval system.
- To identify user needs by collecting data on the subjects and topics that are most searched for or clicked on.
- Effective tracking enables the development and improvement of the user interface and software by analysing user behaviour and navigation patterns.
- It identifies trends and changes in user behaviour and preferences, and ensures effective decision-making regarding the changes to be made in the nature and content of the site.
- Notes the geographical concentration of requests received by the Web site, indicating the shifts in demographic of the target audience.
- Provides information to site owners regarding the number of and location of new visitors.

Over the years, webmasters—the guys who create and maintain Web sites—have realised the above benefits of calculating statistics about Web site usage. Dozens of software products are now available in the market—like NetTracker, SurfSpy, WebTrends, TheCounter.com, etc.—which do the job of tracking for a Web site. However, the utility value of such software is limited and some-

times grossly exaggerated.

2.2.3 Statistics generators

There are basically two kinds of programs that “generate”—or survey and record—statistics of traffic to a Web site: the ones which monitor sites “dynamically” and the ones that study the log files of Web servers.

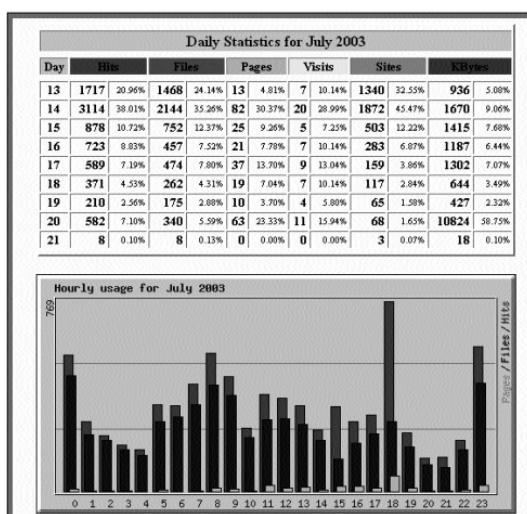
1. Dynamic Monitoring

The monitoring of number of visitors to a Web site in real-time is known as Dynamic Monitoring. Usually, a Web site counter shows the number of hits of the page, and this number is updated every time a new visitor accesses the site. This usually takes the form of a “counter,” or an image (an advertising banner, for example) that is placed on the main page of the site. When your browser requests the image, a program is executed (usually a CGI script or server-side application) which records the IP address from which your request is coming, and then returns the image to the same address. This is great for identifying the different IP addresses that have visited the site.

Usually, a hit counter is a good option for sites that don't have a Web server of their own, and therefore cannot access the more reliable server log files.

2. Analysis of Web logs

Every time a Web site is accessed, a request is actually made to a server for that particular site to be accessed.



A Web tracking software typically produces a fancy report like this one

Each request is recorded in a standard log file format. Being standardised, log files are highly accurate criteria by which the popularity of a Web site could be established. Some of these log files are:

Access_log: The principal log file that records filenames, IP addresses, date, time, and other relevant data.

Referrer_log: Stores the URLs of the sites from which the user was referred to this site.

Error_log: Requests for files and error messages are entered into this file.

Perhaps a really simple way to find how many requests a server received is to count the number of lines in the log, as each log entry represents a request. Besides, each log line shows a request for a specific URL that comes from a definite IP address. This could make it even easier to trace how many requests a particular page on the Web site got and how many came from a single source.

2.2.4 Felling Logs

The way to analyse Web site traffic is to be able to understand what exactly a typical log entry in a Web server actually means. A typical log entry would look something like this:

```
123.123.45.13 - [15/Jan/2008:11:30:22 -0400] "GET/stuff.html  
HTTP/1.0" 200 16716
```

The data in the Access _log file above is made up of several different fields. A file in the Common Log Format typically has eight fields of information for each HTML or graphics file served. Going from left to right, the fields that are listed in the log are:

1. Address field (123.123.45.13): This is obviously an IP address. The address field records the domain name or IP address from which the request has been made. So if you are asking for the Yahoo! home page, the Web server would make a note of *your* IP address (not the IP address of the Yahoo! site) as the first item of data in the access log. Often, system administrators choose to store only the IP address and not the domain names, because that would require

more bandwidth.

The address field is invaluable; it helps tabulate exactly how many different individual user requests have been received by a site. By counting how many different IP addresses have been logged, log analysis programs can determine how many different sessions have occurred and how many times a single source has requested access to the site or any one of its pages.

2. **ID field (-):** This field is usually not used, for security reasons.
3. **AuthUser field (-):** This field would show up only when the log registers that the user was required to type in a username and a password in order to see the page requested. This field usually reveals the full name of the user who owns the account that is making the HTTP request. Most remote sites don't give out this information for security reasons. If this field is disabled by the host, you see a dash (-) as you do here.
4. **Date and time field (15/Jan/2008:11:30:22 -0400):** The date and time when the server sent the requested material into the Internet. This field is essential to calculate the time elapsed between two page displays. This, then, can be used to sum up how long a particular user viewed a particular page. The log notes the details in this order:
 - Date (15/Jan/2008)
 - Time—hours, minutes, seconds (11:30:22)
 - GMT offset, meaning the difference in time between the time noted here and the GMT. In this case, the time when the data was served to the user is four hours earlier than GMT.
5. **Method field (GET):** Gives information about the request itself. For all HTTP requests (the www page requests), this field will always start with the GET command. Most Web pages are served using the GET command in HTTP (which is the language in which Web browsers and Web servers talk to each other). Sometimes, the GET is replaced by POST. The difference here is not so relevant to the guy who's interested in Web statistics, hits and misses, etc. The reason is that POST basically indicates that the forms pass their data "invisibly" to the server, while forms that use the GET method pass the data directly as a part of the URL to the server.
6. **Filename field (/stuff/stuff.html HTTP/1.0):** This one is self-evident, in the sense that it shows the path and the filename of the

requested site—in this case, the client has requested the resource /stuff.html, which means that the file asked for is an HTML file and not a file in the.gif or.tif format. The second bit of information that we get from the filename field is that the client used the protocol HTTP/1.0 to access the file.

Note that the Method field and filename field information are always presented in double quotes.

7. Status field (200): The second-to-last number at the end of the line represents the response code of the request. This response code is basically a code that is a part of the HTTP protocol, and which shows how exactly a Web server handled a request. The status field, therefore indicates whether this particular request was successfully delivered by the server. Here, the code is 200, which indicates successful delivery of the file. This field is useful for detecting if a particular file has remained undelivered due to errors. Some of the typical status codes you'd see in the log are...

Code	What it means
200	Successful delivery of the file
404	File not found
302	Redirect to another file
401	Password required
500	Server error
400	Bad request

You're probably very used to seeing the message “404 not found,” which appears when your request doesn't find the page you requested.

8. Size field: The field which specifies, in bytes, the size of the file that has been sent to the client. Typically, Web server software report the size of the file served. In this case, it is 16,716 bytes.

2.2.4 Analysing Log Files

So now that you know what log files actually represent, it's easy to understand what the owners of Web sites could do with this data: they can learn about their audience and their information needs. By following the path—the links—that people take while going

through a site, visitors' behaviour can be studied and patterns can be traced, which could assist in making the site more attractive and user-friendly.

Web analysers can use all the above fields (and several others which we haven't mentioned) to count the number of hits, unique visitors, repeated visits, user preferences, geographical location of requests, etc. And none of this needs to be done manually, counting each and every line in a server's log and finding similarities or differences. Software do all the counting that needs to be done, and even produce intricate charts and graphs to visually assist one in the understanding of the data so calculated.

However, it's not all a cakewalk. Several inaccuracies are possible, and Web sites would grossly misconstrue the behaviour of their target users if they considered the analysis of Web logs to be foolproof.

2.2.5 Statistics That Don't Stick

There are a lot of things that can't possibly be accurately calculated by any analysis program, even though commercially-available ones would sometimes have you believe otherwise.

1. Same IP address, different user

There is no way in which a Web site analyser can identify how many real people actually visited or requested a particular Web page. There are several possible reasons for this.

- Typically, ISPs assign a different IP address each time a user connects to the Internet. So it is quite possible that once you log off your ISP server, another user is assigned the same IP address that you were operating from and accesses the same Web site that you did. A Web server log would simply note that the request has come from the same source and consider that either the visitor is non-unique (the same visitor has been navigating through Web pages on the same site for a longer session).
- Online services (like AOL and MSN) often allow several hundred users to access the Internet through a relatively small number of gateways. Each gateway has its own IP address. Therefore, several individual PCs would be recorded in the log as a single client.
- Corporate users allow employees to connect to the Net through proxy servers. Each PC in the sub-network of the organisation

uses the same IP address to access the Web—that is, the IP address of the proxy server.

- Universities, libraries, and cyber-cafes allow lots of people to access the Web using the same computer, the same IP address, which can be interpreted as the same user making repeated visits to the site.
- “Spider” or “robot” software are Internet robots (used by search engines) that explore the Web and collect Web page addresses based on content found at those pages. Often, a page view is just a spider crawling through, but is reported as an individual client with a separate IP address.

2. Same Web site, different servers

Sometimes Web sites are so large that they need to spread out their pages on multiple servers to enable easy access to the large number of requests they receive at the same time. Duration of a page view is therefore only really assessable on individual servers. In such cases, the quantum of time spent on a Web site would not be visible on every individual server that holds a part of the content of the Web site.

3. Size Matters

Some Web server programs always record the exact, actual size of the file that is being sent. In such a case, if a particular user stops downloading large graphics files mid-way through the process, the server still reports the file size, instead of the actual amount of data served.

4. The Referrer Field

Often, a log file shows the URL of the page that contains the hyperlink you used to access the current page. This enables the maintainers of the Web site to tell which Web site referred the user there. The referrer field reveals which page the visitor was on when he clicked to come to this page. Usually this will mean that the page had a link to that particular site. But sometimes this is simply the page the user was looking at when they typed in your address into their browser, or clicked on your address in some other software such as a newsreader or an e-mail client. Using the

referrer field as a reliable source of data for where the traffic to your site is coming from would therefore be a mistake.

5. Counter Discourter

Using the primitive page counter system has its own problems. This is usually a dynamic bit included in a Web page that increments a counter that is displayed each time a page is requested. This is usually a part of the Web site as a different “image” file, which is loaded every time you access the page. However, if the user has

turned off the image display option in his browser, the images would not be displayed at all, leading to one count less in the counter. If a user just reloads a page, the counter registers another page view, not figuring that it’s the same user viewing the same page again.



A hit counter could look something like this...

Welcome to the Page Counter page.

You're visitor number: 00001

....or like this

6. Spooky Cookies

Cookies may get into your system, but if you change the settings in your browser to block all Web sites from installing cookies on your computer, a Web site would be unable to identify you as an earlier user, and log you as a new one.

7. Tousled Browsers

The log, on a server, could also contain a “browser field” (also called the “User Agent” identifier), which is usually a piece of text that specifies what software was used to access the site. This, however, can be faked; Opera 8, for example, could fake being Internet Explorer to open MSN Web sites and other related Web pages which deliberately set up different styles that would not be compatible with non-IE browsers.

2.2.5 To Sum It Up

So there's only so much that Web statistics can tell us for sure:

- How many times a Web site is accessed from the same IP address in a particular duration of time
- The number of IP addresses that request data from a Web site
- When exactly each request is made
- The most popular URLs on a particular site
- The particular day of the week or the time of day that the server remains most active
- The geographical location of the servers that request access the maximum number of times.
- The kind of files that are most requested from a particular IP address (JPEG, HTML, etc.)

Ideally, calculating visits on a per-page visit is more reliable. The better log-analysis tools available in the market can rank pages on a Web site, showing which are most popular and which pages are rarely or never visited, enabling webmasters to modify the content accordingly.

Section II

Software

Web Browsers

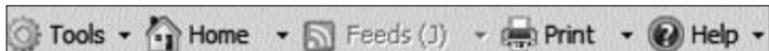


The early adopters of the Web quickly realised that getting computers connected to the Internet was not enough; they needed software that would allow the guest computer to browse the information hosted on those connected computers. This thought process lead to the formation of the Web browser. Let's take a closer look at all these browsers in the next few pages, starting with the most familiar.

3.1 Internet Explorer

Microsoft has bundled every new release of Internet Explorer (IE) with Windows. The most important aspect of IE is that since it has a market share of approximately 75 per cent, all Web pages are designed to work with it. Apart from that, since IE is bundled with the OS, there is no need to use an alternative browser.

IE 7 has gone through a complete redesign and now supports the opening of multiple pages in one window—also called tabs. Apart from that, the update plugs many security issues with IE 6 and earlier versions. The new design also incorporates an anti-phishing filter and better security settings.



The default icons on the Command Bar

The new design needs some getting used to initially but once you get used to it, you will find it to be more intuitive than IE6. To customise it, right-click anywhere on the toolbar. You can now choose to remove the Menu bar, Links or the Status Bar. You can also reduce the size of the icons by un-checking Use Large Icons. The Customize Command Bar option on this menu allows you to customise the icons visible to the right of the Toolbar.

Just above the Command Bar, you have the search box from where you can search your queries. To access this box quickly, use [Ctrl] + [E]. By default, this box uses Windows Live as its search engine. You can, however, change that to any other search engine—Google, for example. Here's how:

1. Click the arrow on the magnifying glass icon next to the Search Box.
2. Now, click Find More Providers...
3. In the page that opens up, click on the search engine you would like to use and select Install on the next screen.

If you want to use this search engine by default every time you run a search then click the arrow on the magnifying glass again and select Change Search Defaults. Highlight your search engine and select Set Default. You can also remove the other providers if you don't plan on using them in the future.



You can also quickly bookmark a page by clicking on the icon to the left of the first tab and clicking Add to Favourites. Alternatively, you can use [Ctrl] + [D].

If you want to browse a page or RSS feed you have bookmarked, Click on the star icon on the toolbar and select the page accordingly.

One of the biggest changes seems to come from the focus of allowing third-party developers to build add-ons to increase the functionality of the browser. These add-ons start every time IE is started and can be called upon to do various tasks such as managing passwords to checking your spellings as you type. Let's take a look at some of the best add-ons available for IE:

3.1.1 IE7 Pro

No IE7 add-on list is complete without this one. When the browser was first launched, it disappointed quite a few by providing too few features too late, after all, alternate browsers like Opera and Firefox had been around for a while and offered better functionality. Then IE7 came along—an add-on that carried enough oomph to make using the browser worthwhile.

What does it do, you ask? Plenty! It has a spell checker, a better ad-blocker than that available in the browser, an auto form filler, support for user scripts and mouse gestures, among other things. After installing it, the first thing you should do is to switch on Tab History. This will allow you to bring back any tab you may have accidentally closed in the session. Mouse Gestures allow you control functions like going back and forth on a Web site, opening or closing new tabs or automatically scrolling a Web page without having to press any buttons on the screen. The ad filter is efficient, and allows you to create your own filters to block domains. The Search option allows you to search from the address bar, just like Opera. To search with Google, enter "g" before your search query. There is also the ability to add custom search providers from the Search menu.

The Online Service menu lets you synchronise your bookmarks automatically with IE7Pro's storage servers online. When you select this option for the first time, you will be prompted to make an account. After that, it will start to sync your bookmarks automatically online. You can also choose to define your own custom shortcut

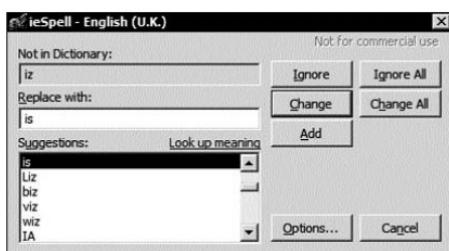
key combinations or modify existing combinations to suit your taste. The Web site www.iescripts.org provides IE7Pro scripts, styles, and plugins, which you can use to customise IE even further.

Though the User Scripts option carries a lot of potential—just like GreeseMonkey—on Firefox, they are buggy to an extent and tend to crash the browser more frequently than normal.

You can find this plugin on this month's DVD.

3.1.2 ieSpell

What if you want only a spell checker without the other features of IE7 Pro? ieSpell is what you need. This plugin is designed to scan text from any field on the page. So if you are composing an e-mail or filling out a form on a Web page



ieSpell can even look up meanings of words

and want to check your spelling, you click Tools (on the toolbar) > ieSpell. This will then bring up any spelling mistakes and give you the relevant suggestions, just like in Microsoft Word.

The options menu allows you to ignore uppercase, mixed case words and words with numbers in them. You can also ignore URLs and e-mail addresses. It is best to enable the latter option while leaving the rest unchecked. ieSpell also allows you to add custom dictionaries like those for scientific names.

Let's add a custom dictionary in four simple steps:

1. First download the custom dictionaries to your computer, if they are in TXT format, rename them to DIC.
2. Go to Tools > ieSpell Options > Language and Dictionaries > New.
3. Navigate to where the dictionary is stored on your computer and select it.
4. Select the dictionary to enable it.

You can also look up meanings of words using ieSpell, to do so, just highlight the word and go to Tools > ieSpell. The default dictionary used for checking meanings is Merriam Webster. You can



change this dictionary to Wikipedia, Yahoo! Education Dictionary, Dictionary.com or Cambridge Advanced Learner's Dictionary. You can add more online dictionaries that list meanings too.

Let's add a custom meanings dictionary to the list now, in three steps.

1. A good alternative online dictionary to the ones mentioned above is <http://www.thefreedictionary.com>, so let's add that.
2. Select Tools > ieSpell Options > Others > Edit List > Add.
3. Now, enter the URL <http://www.thefreedictionary.com/%s> in the URL field and name it as Free Dictionary in the Dictionary Name field.

If you want to use this plugin for commercial purposes, you need to buy it from www.ieSpell.com for US \$15 (Rs 616); but get for free for personal use from <http://www.iespell.com/download.php>.

3.1.3 The ParentalControl Bar

If you have kids at home and worry about inappropriate Web sites popping up while they are browsing, look no further than the ParentalControl Bar.

This plugin is very easy to use and offers two modes on installation—Parent and Child. Quite obviously, the Parent mode allows for unrestricted access while the child mode allows for limited access to the Internet. When you first install the ParentControl Bar, it updates the program and prompts you to create a password-protected parent account. Once that is done, you are allowed to login, change the settings of your child account, and choose what you want to block.

Once you've chosen the appropriate settings for the ParentalControl Bar, you can then change from parent to child mode to test the settings. When browsing in the parent mode, you can add sites to the safe list, which can then be viewed by your child when he / she is browsing the Web. You can alter the settings at any time by clicking on Change Parental Settings on the toolbar. The KidTracker option on the toolbar allows you to view all the sites that have been browsed in the child mode.

If you ever forget your ParentalControl password, all you have to do is go to ParentalControl Bar > Forgot Password. The browser will e-mail you the password on the address you specified

during registration. If you have any grievance or suggestion to make the add-on better, you can click on ParentalControl Bar > Report a Concern and fill in the details on the online form. You can download ParentalControl from <http://www.parentalcontrolbar.org>.

3.1.4 RealTime Cookie And Cache Cleaner

The RealTime Cookie and Cache Cleaner (RTC3) is not strictly an IE add-on. It is a program that monitors the cache levels and cookies on the computer, and automatically cleans them when it detects any cookies or when the cache exceeds a certain size. The software can be configured to start with the computer, and resides in the system tray.

There are four sections to the software—Cookies, Cache, Cookie Jar and Index.dat. The Cookies section will monitor cookies created on the computer and delete them as and when they are created. You can specify what cookies you want to keep and delete in the Cookie Jar section. The Cache section monitors the Temp and Temporary Internet Files folder on the computer and once these folders reach a certain size limit, their contents are automatically erased. The Index.dat section monitors all index.dat files on the computer. Microsoft Office and Windows create these files to store details of recently opened files. These files are mostly unnecessary and if not cleaned periodically can take up unnecessary space on the computer.

RTC3 does not take too many system resources to run in the background either. If you don't want it to continuously scan in the background, you can disable that under Options. You can also choose if you want RTC3 to load on startup or have it clean the history, system and Internet cache and history every time the computer boots. You can also set the scan or clean interval from here. It is better to set the level to 60 seconds or more on slower machines. You can download it from <http://www.kleinsoft.co.za/>



Cookies being tracked and deleted by RTC3.

3.1.5 Shortcuts

General Shortcuts	
Turn Full Screen Mode on or off	[F11]
Cycle through the Address Bar, Refresh button, Search Box, and items on a web page	[TAB]
Find a word or phrase on a page	[Ctrl] + [F]
Open the current webpage in a new window	[Ctrl] + [N]
Print the page	[Ctrl] + [P]
Select all items on the page	[Ctrl] + [A]
Zoom in	[Ctrl] + [+]
Zoom out	[Ctrl] + [-]
Zoom to 100%	[Ctrl] + [O]
Navigation shortcuts	
Go to home page	[Alt] + [Home]
Go backward	[Alt] + [Left Arrow]
Go forward	[Alt] + [Right Arrow]
Refresh page	[F5]
Refresh page and the cache	[Ctrl] + [F5]
Stop downloading page	[Esc]
Favourites Centre shortcuts	
Open Favourites	[Ctrl] + [I]
Open Favourites in pinned mode	[Ctrl] + [Shift] + [U]
Organize Favourites	[Ctrl] + [B]
Add current page to Favourites	[Ctrl] + [D]
Open Feeds	[Ctrl] + [J]
Open Feeds in pinned mode	[Ctrl] + [Shift] + [J]
Open History	[Ctrl] + [H]
Open History in pinned mode	[Ctrl] + [Shift] + [H]
Tab shortcuts	
Open link in new background tab	[Ctrl] + left click or middle click
Open link in new foreground tab	[Ctrl] + [Shift] + left click or [Ctrl] + [Shift] + middle click
Close tab (closes window if only one tab is open)	[Ctrl] + [W] or [Ctrl] + [F4]
Open Quick Tab view	[Ctrl] + [Q]
Open new tab	[Ctrl] + [T]
View list of open tabs	[Ctrl] + [Shift] + [Q]
Switch to next tab	[Ctrl] + [Tab]
Switch to previous tab	[Ctrl] + [Shift] + [Tab]
Address Bar shortcuts	
Select the Address Bar	[Alt] + [D]
Add “http://www.” to the beginning and “.com” to the end of text in Address Bar	[Ctrl] + [Enter]
Add “http://www.” to the beginning and the website address suffix you have specified to the end of text in the Address Bar*	[Ctrl] + [Shift] + [Enter]
Open the website address that is typed in the Address Bar in new tab	[Alt] + [Enter]
View list of previously typed addresses	[F4]
Search Bar shortcuts	
Select the Search Bar	[Ctrl] + [E]
View list of search providers	[Ctrl] + [Down Arrow]
Open search results in new tab	[Alt] + [Enter]

3.2 Maxthon

This browser was created for the IE 5.5 days, when newer browsers with better functionality were appearing on the horizon and Microsoft couldn't care less about updating IE. Even when they finally did, it included functionality that was—by then—almost commonplace in all the alternative browsers. This is when Maxthon was released.

This cannot be called a browser in itself because it borrows the Trident rendering engine used by IE to display Web pages and is more like a GUI upgrade for IE. When it was launched, its tab functionality, password manager, and ability to use IE favourites received quite a positive response from Internet users. Since then, IE 7 incorporates nearly all the features of Maxthon. In spite of this, its fan base remains, albeit diminishing by the day. Maxthon recently released a newer version of their browser, which added more features.

When installing Maxthon 2, you are asked if you want to create an account to store your settings for different users or use one profile for everyone. A few features are not available when using the single profile setup and therefore we choose to create an account. You can choose to automatically login to your account (quite pointless; if you want to share your settings, you'd rather not make that account in the first place!) or enter a password every time you start Maxthon.

Once you are logged in, you can choose to synchronise your bookmarks with an online database. This way you can use any Maxthon browser in the world and log into your account to find all your bookmarks there. Apart from this nifty feature, the browser has an inbuilt Web accelerator called “Super Maxthon Smart Acceleration”, which will monitor the Web sites you visit, and cache its images and layout so that they load faster.

Maxthon also has an inbuilt RSS Reader, support for tabs, it can restore the tabs when the browser crashes, has an undo option for when you accidentally close a tab, can block ads, flash, ActiveX and Java applets. It also supports mouse gestures and an automatic update service.

The unique functionality of this browser is in a feature called Groups, which allows you to store bookmarks under a folder. When



you click on a group, all the bookmarks stored under it open in separate tabs at once. You also have the option of creating aliases for long URLs that you visit frequently—type ‘yahoo’ instead of ‘mail.yahoo.com’ to visit your Yahoo! Mail login page, for example.

URL Key	
URL Key settings will override your settings in the Custom Shortcut Keys section. You can input multiple URLs, one per line.	
Shortcut Key	URL
F1	www.gmail.com www.yahoo.com www.hotmail.com
Add Item	

An example of the URLs associated with the F1 key in Maxthon.

Let's create this alias in three easy steps now:

1. In the browser, go to Tools > Maxthon Setup Centre
2. In the new tab that opens click Navigation and scroll to URL Alias
3. You can now enter the alias name and the URL.

Apart from this functionality, you can also use the function keys to open a group of URLs, all you have to do is go to Tools > Maxthon Setup Centre > Navigation and scroll to URL Key. You can now select the function key you want to use along and enter the URL addresses you want to associate with that key.

The browser also has a small Notepad-like utility called Simple Collector that can be used to store URLs, text or images from sites, or to write JavaScript scripts and run them on the open Web page.

Another queer option in the browser is Split View. What this does is divide the browser into two columns, each with its own set of tabs. This is most useful when you're using a dual monitor setup. You can select a column by clicking into it and drag-drop tabs from one column to the other.



The Split View functionality in Maxthon

3.2.1 The Skins

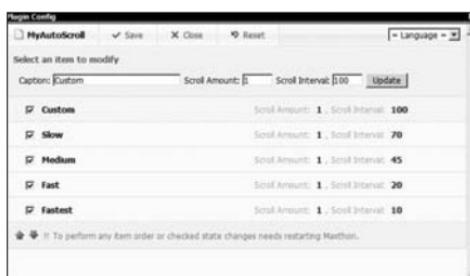
Unlike IE, Maxthon supports skins, which allow you to change the GUI of the browser. It comes with the Maxthon 2.0 skin preinstalled along with the Maxthon Classic Skin. There are tons of skins available online and to access them, you can go to <http://addons.maxthon.com> and select MX2 Skins. To learn how make a skin, you can go to <http://tinyurl.com/3dgrau>.

3.2.2 Add-Ons

Like IE, you can add many add-ons to Maxthon too. These can be found at <http://addons.maxthon.com>. The add-ons are divided into various categories on the Web site like Browsing Tools, Entertainment, and Mail/Search among others.

Better BugMeNot (<http://tinyurl.com/2w29pt>) is an add-on that searches for the usernames and passwords from www.BugMeNot.com for Web sites you don't want to register on. After installing the add-on, you go to the login page of the Web site you want to access and click the BugMeNot button on the toolbar. The add-on will fetch the username and password and insert it in the relevant fields on the page, saving you quite a lot of time.

Another add-on worth mentioning is My Auto Scroll (<http://tinyurl.com/2h2257>). This add-on will automatically scroll any Web page for you when you are browsing. You can define



The Configuration settings of MyAutoScroll

the speed of the scroll in the Configuration menu of the add-on.

VideoDownloader (<http://tinyurl.com/2yr6he>) is another handy add-on to use. As you might have already guessed, this one allows you to download videos directly from www.YouTube.com. After installing the add-on, you just have to click on the VideoDownloader button on the toolbar when you want to download a video from YouTube.com. This add-on also works with other online video sites like Google Video, Metacafe and iFilm. The videos are downloaded as FLV files which

can be viewed in VLC Media player or Media Player Classic—you can find them in the Essentials section of any of our CDs.

3.2.3 Tips And Tricks

When browsing Flash animations and video on Maxthon, you'll run into a few scripting errors. This is because of the limitation imposed by Microsoft on third party browsers using the Trident engine. There is, however, a solution to this:

1. Go to Start > Run and enter Regedit.
2. In the new window that opens, navigate to HKEY_LOCAL_MACHINE\Software\Microsoft\Internet Explorer\MAIN\FeatureControl\FEATURE_BLOCK_LMZ_SCRIPT. If it doesn't exist, you can create it as a new Key
3. Right-click in the right hand column and add a Dword with the name Maxthon.exe with a value of 0.

You can also selectively enable the Maxthon Smart Acceleration option for Web sites you visit often. If you want to browse quickly through open tabs, use the mouse wheel while the cursor is on the tabs. Alternatively, you can also use [Ctrl] + [Tab]. When the Split View is activated, middle clicking a link will always open it in the split that has focus.

There is also a “CPU Saving Mode” tucked away under Tools. This has two options—Standard and Full. The Standard option allows you to disable flash and JavaScript in tabs that are not in focus and the Full function allows you to disable those elements completely in all the tabs of the browser. If you want to disable images on the page, you can do that too. Go to Tools > Content Central from where you can disable not only images but also scripts, flash movies, sounds and ActiveX objects on a Web page.

If you want all the features of Firefox with the Web page rendering compatibility of IE, Maxthon is your best bet.

3.3 Firefox

A lot has been written about Firefox (FF) and over the last few years, it has been the browser to take the fight back to the IE camp with its slogan “Take back the Web”. If you were living in the Himalayas (we’d like to do that someday) for the last few years then you would not have noticed that the browser market has steadily started getting more and more fragmented—particularly in the last two years. With the ‘Spread Firefox’ campaign getting more and more popular around 2005, more users realised that it was able to browse faster than IE, was more Web standard-compliant, and boasted of better security at the time. Add to that tabbed browsing and the ability to add more functionality, and quite a few Web surfers shifted loyalty to FF in droves.

FF is a free and open source browser, which is available for the Windows, Linux, and Mac platform and relies on the support of the Web community to build new features for the browser and test them. The ideology behind the browser, according to its developers, is to deliver the “best possible browsing experience to the widest possible set of people.” The latest version of FF is version three, which is currently in beta. The FF browser allows tabbed browsing, an in-line spell checker, bookmarking and a download manager among other features.

3.3.1 Add-Ons

Like IE and Maxthon, FF also allows users to add functionality to the browser by adding add-ons. Its biggest advantage however lies in the fact that FF has the more number of add-ons than IE and Maxthon combined. As of January 2008, it is estimated to have over three thousand add-ons. In FF, add-ons are divided in three distinctive groups: **Plugins:** These help the browser perform specific tasks—viewing certain graphic formats, playing certain multimedia file types, etc.

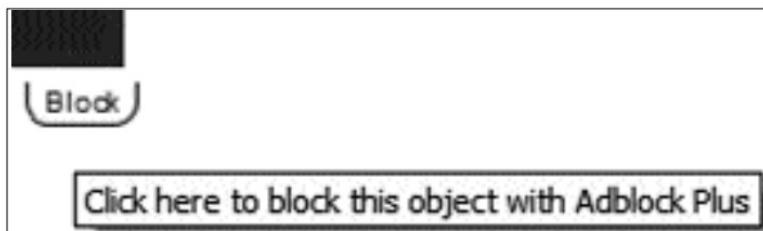
Extensions: These are small add-ons that add new functionality; this can range from a simple toolbar button to a completely new feature in the browser.

Themes: These allow you to change the UI of the browser, like in Maxthon.

When you visit a page that hosts files not compatible with FF, it will request you to download the relevant plugin (if available). All the most important plugins can be downloaded from <http://tinyurl.com/29sbvj>. If you surf a variety of Web sites every day, it is best to download all the plugins from the above page.

As mentioned above, there are many extensions available for FF. However, as you keep installing extensions, the browser uses more RAM and surfing the Web can get slower. Also, the chances of security vulnerabilities increase with the number of extensions you install—there is a thin line to tread between functionality and usability in FF. Having said that, there are some great extensions available for FF, here are some of them.

AdBlock Plus: This is a plugin that well... blocks ads. What is so special about it is that AdBlock Plus subscribes to the EasyList Filter, which maintains a list of URLs which is updated by users and blocks them. The plugin will also allow you to block any image or script string on the Web page. If many users have blocked the same URL, it gets updated on the list online, thereby allowing it to block ads more effectively. Apart from that, there is also a white-list function, which will allow ads to display in the future. The EasyList Filter is updated every few days so there is very little chance of you encountering an ad. Even if you do, you can block it by right clicking on the ad and selecting Block Ad.



The easiest way to block an image.

Greasemonkey: This plugin allows users to install scripts that can make changes to the layout of most HTML Web pages on the fly. Since the script runs in the background every time the page loads, the change is permanent for the user. If you want to return to the previous state then you have to stop running the script from the

Greasemonkey extension. As an example, the Enhance Gmail Greasemonkey script is used to enhance the functionality of Gmail. Other users have written greasemonkey scripts that auto fill forms, remove ads, pop-ups and Google text ads from Web sites. There is also a script created for the Indian Railways Web site (<http://irctc.co.in>) which modifies the interface to show the availability of seats in the train while booking your ticket. You can find it at <http://tinyurl.com/2wp9kb>.

To install a Greasemonkey script, you have to install the extension first from <http://tinyurl.com/ynvmda>. Then you can continue to <http://userscripts.org> to search and download the scripts you want to use. It is best to browse the latter site first. There are quite a few plugins that you may not know existed, e.g. adding the Gtalk chat application in Orkut etc.

Tab Mix Plus: Here is one of the most widely used extensions in FF. With this, you have a dozen options ranging from tab duplication, merging and focusing. Apart from that, this extension also has its own robust session manager with a lot of extra functionality.

There are quite a few features, which are unique to Tab Mix Plus, one of them being the ability to open JavaScript popups in a new tab or block them from opening altogether.

Tabs can also be configured to open independently from either the search bar, bookmarks menu or the address bar—for example, you can configure a new tab to open up every time you visit an address after entering it from the address bar. You can also specify if you want a new tab to open with a blank page, the home page, duplicate the currently focused tab, or any user-specified URL. If you thought that was too much, the extension also allows you to choose between twenty-five options such as Undo Close Tab and Lock Tab in the Tab Context menu.

Tab Mix Plus also has a full-fledged Session Manager, which records all the tabs open at any given time, and provides various



Some of the settings of Tab Mix Plus.



options for saving and re-opening those tabs. If the browser crashes, its Advanced Recovery option will bring back all the tabs that were open at the time of the crash. You can also set mouse click combinations to perform various tasks on the tab bar and browser toolbar.

The plugin is by far the best way to manage tabs and makes browsing that much easier. You can download the extension from <http://tinyurl.com/2au5zj>.

Cooliris Previews: Haven't we all clicked on links on a Web page that took us to a completely unrelated page? Isn't it frustrating when that happens? Well you can avoid that hassle now thanks to the Cooliris Previews extension.

When installed, this extension will allow you to preview the link you are going to click on even before you have clicked it! It works like this: When you hover your mouse over any link, a small magnifying glass will appear beside it. When you place your mouse over the icon, a small window opens up showing you the contents of the link you were going to click on.

That's not all, from the window that opens up, you have an option to e-mail the page to a friend or open the current link in a new tab. You can also increase or decrease the text size in the window or go back and forth between links on the page from there. Another functionality called Stacks allows you to bookmark a link temporarily on the side of the page.



Various Tab Context Menus with Tab Mix Plus



An example of the stacks and previews functionality in operation in Cooliris Previews

There is another nifty use of this extension. When you highlight a block of text on a Web page and right-click in the selection, you get the option to search that text right from the context menu!

The options menu of the extension allows you to define how long to wait before the preview windows appears on screen. There is another setting which allows you to define which site you want to use to search the block of text from the context menu. This can be configured to any search engine of your choice. There already a few pre-configured sites like Yahoo!, Flickr, eBay, Google News, etc. which you can use. You can download Cooliris Previews from <http://tinyurl.com/2xxouc>.

UnPlug: Like the VideoDownloader add-on of Maxthon, UnPlug is a similar extension that will download hidden content from Web pages. It works slightly differently from VideoDownloader, though. When you reach a page which hosts the embedded content—that may be video, flash animation, an audio stream or a java applet, you click on the UnPlug button on the toolbar. The extension will then analyse the page for the embedded content and display the relevant links to you. All you have to do is click the ‘save’ button next those links to download the content.

The best part of UnPlug is that it doesn’t need to connect to any external sites to download the content. If you have other download assisting extensions installed Unplug will integrate with them, e.g. *Download Them All* and *NoScript*. You can download UnPlug from <http://tinyurl.com/2rnc22>.

ChatZilla: Another very popular extension used is ChatZilla. This helps users to chat on IRC from within the Firefox browser itself. After installing it, go to Tools > Chatzilla; a new window will open up. Since ChatZilla is a barebones IRC Client, you have to be proficient with IRC commands to be able to connect to the right server and channel without having to rely on a helpful GUI. Having said that, there is a lot of helpful documentation available online for IRC commands, a good place to start would be <http://tinyurl.com/2nfxrw>. You can download ChatZilla at <http://tinyurl.com/227n3v>.

All-in-One Sidebar: This extension acts as a quick menu from where you can control downloads, the extensions installed on your computer and the browser history, among other things.

When installed, the All-in-One Sidebar is located on the left hand edge of the screen. From here, you can open various sections such as Bookmarks, History, Downloads, Add-Ons, MultiPanel, Page Info, Page Source and Error Console. Each of these sections performs a specific task—for example, the MultiPanel section allows you to browse other Web sites in the smaller panel while the Page Info section will show you the detailed characteristics of the Web page you are viewing. The Error Console displays all the errors in the code of the Web page. These are useful to determine problems with the layout quickly and can be quite helpful for site designers or novices learning to code Web pages. Download All-in-One Sidebar from <http://tinyurl.com/2k54f6>.

3.3.2 Tips And Tricks

Though FF is pretty efficient on its own, there are various tweaks you can carry out to improve the layout and the speed of the browser. Here are some tips and tricks to help you realise that dream:

1. To increase more screen space for the Web page, make all your icons small. To do so, right-click anywhere on the toolbar and select Customise. Now, right at the end of the window, put a tick mark next to Use Small Icons.
2. It is possible to search with different search engines right from your address bar window. All you have to do is go to your favourite search engine, right click on the search field and select “Add a keyword for this page”. So, if your search engine is Google, you may add ‘g’ as my keyword. Now, when you want to search using Google, all you have to do is enter “g [Space] [search query]” (without the quotes) in the address bar.

If you want to customise FF further then you have to use the about:config tool. To access this, enter “about:config” in the address bar and press [Enter]. This is a place where all the layout and page rendering settings—not normally accessible from the menus—are stored. You can change many of the settings to improve the way FF works.

We must mention here that using this tool without knowing what you are doing could cause FF to stop working. It is therefore important to follow these steps carefully.

If you don't like the close buttons on all tabs and accidentally press them when switching between tabs, you can turn off this functionality by entering 'about:config' in the address bar and searching for `browser.tabs.closeButtons`. Change the value to 3 to remove the close buttons from the tabs; 0 will show the close button only on the active tab.

The way FF speeds up your browsing is by automatically downloading content from links on page it thinks you will click on—the first search result, for example. This uses up a lot of your bandwidth unnecessarily. To stop this, search for `network.prefetch-next` in the about:config tool and set its value to false.

The one of the biggest annoyances in FF is that the in-line spell checker does not work in text fields and Web forms. You can enable this by search for `layout.spellcheckDefault` and setting the value to 2. While you are at it, you can also search for `browser.urlbar.hideGoButton` and set the value to true. Who uses the 'Go' button on the toolbar anyways?

The browser waits 250 milliseconds before it renders any new content. To render pages quicker and remove the wait time, go to about:config, right click in the window and select New > Integer and call it `nglayout.initialpaint.delay`. Now set its value to 0. FF will not wait anymore before it renders new content on the page.

Have you searched from the Search Bar only to be annoyed when the search results overwrite the contents of the existing open tab? You can now automatically open results in a new tab by going to `browser.search.openintab` and setting it to true.

Another annoying functionality of FF is to open bookmarks in existing tabs, if you want to open those on a new tab too, search for `browser.tabs.loadFolderAndReplace` and set it to false.

If you are in the habit of opening many tabs, you would've noticed that they don't wrap around the toolbar, this way it is tough to know all the tabs that are open at any given point of time. One way to sort this is to make the size of the tabs smaller. To do this, search for `browser.tabs.tabMinWidth` and reduce the value to 50. This will shrink the tab size, making more tabs fit on the tab bar. We must mention though that doing so will also make it difficult to read the title of the tab.

If you browse multiple Web sites at once, it is better to increase the number of connections FF makes to servers. This will increase your browsing speed to an extent. Search for network.http.max-connections and set the value to 32. Then, search for network.http.max-connections-per-server and set the value to 8. Now, search for network.http.request.max-start-delay and set its value to 0. You can also change network.http.max-persistent-connections-per-server to 8 as some servers will allow up to eight persistent connections per request.

Normally when you send a request to view a page, the browser will wait for a reply from the server. FF contains an experimental browsing acceleration feature called pipelining that sends out multiple requests to the server at once without waiting for one reply after the other. To enable this feature, go to network.http.pipelining and set its value to true and change the value of network.http.pipelining to 16, if you are surfing through a proxy, set network.http.proxy.pipelining to true as well.

The amount of RAM used by FF increases as you keep using it though the day. There are, however, a series of tweaks that should keep that appetite in check. Search for browser.cache.memory.capacity and set the value to 15000 for RAM sizes between 512 MB and 1 GB. For lower sizes, choose 5000. The browser keeps all the images and pages that have been rendered in cache so that if you open any page from History, the load times are quicker. This setting however, drives up the amount of RAM as you access more and more pages. You can disable this option by searching for browser.cache.memory.enable and setting its value to false. This will however reduce the load times of pages accessed from History.

A little known setting in the browser allows you to reduce the amount of memory used by the browser when it is minimised by not closed. To activate this, create a new key by right clicking on the about:config window and selecting New > Boolean and naming it config.trim_on_minimize. Now set the value to true.

Here are some nifty shortcuts you can use to enhance your browsing experience:



Shortcuts

Command	Shortcuts
Help	[F1]
Add Bookmark	[Ctrl] + [D]
Bookmarks Sidebar (toggle)	[Ctrl]+[B] [Ctrl] + [I]
Browsing History (toggle)	[Ctrl] + [H]
Download History	[Ctrl] + [J]
Clear Private Data	[Ctrl]+[Shift]+[Del]
Full Screen (toggle)	[F11]
Increase Text Size	[Ctrl] + [+] [Ctrl] + Scroll down
Decrease Text Size	[Ctrl] + [-] [Ctrl] + Scroll up
Restore Text Size	[Ctrl] + [0]
Print	[Ctrl] + [P]
Reload Page	[F5] [Ctrl] + [R]
Force Reload Page (override cache)	[Ctrl] + [F5] [Ctrl] + [Shift] + [R]
Stop Loading	[Esc]
Save Page As	[Ctrl] + [S]
Page Source	[Alt] + Left-click [Ctrl] + [U]
DOM Inspector	[Ctrl] + [Shift] + [I]
Go Down One Line	[Down Arrow] [Alt] + Scroll down
Go Up One Line	[Up Arrow] [Alt] + Scroll up
Go Down One Screen	[PageDown] [Spacebar]
Go Up One Screen	[PageUp] [Shift] + [Spacebar]
Go to Bottom of Page	[End]
Go to Top of Page	[Home]
Move to Next Frame	[F6]
Move to Previous Frame	[Shift] + [F6]
Move Focus to Next Actionable Item	[Tab]
Move Focus to Previous Actionable Item	[Shift] + [Tab]
Open Home Page	[Alt] + Home
Back one Page	[Alt] + [Left Arrow] [Ctrl] + [Shift] + Scroll down [Backspace]
Forward one Page	[Alt] + [Right Arrow] [Ctrl] + [I] [Shift] + Scroll up [Shift] + [Backspace]
Open Link	[Enter] Left-click
Open Link (in new Background Tab)2	[Ctrl] + [Enter] [Ctrl] + Left-click Middle-click
Open Link (in new Foreground Tab)2	[Ctrl] + [Shift] + Left-click [Shift] + Middle-click



Command	Shortcuts
Open Link in New Window	[Shift] + [Enter]
Save Link Target As	[Shift] + Left-click
Caret Browsing	[Alt] + [Enter]
Open New Tab	[F7]
Close Tab (close Window if only one tab)	[Ctrl] + [T] Double-Click on Tab Bar
Undo Close Tab	[Ctrl] + [W]
Next Tab	[Ctrl] + [F4]
Previous Tab	Middle-click on Tab [Ctrl] + [Shift] + [Tab]
Select Tab [1 to 9]	[Ctrl] + [PageUp]
Open New Window	[Ctrl] + [Shift] + [T]
Close Window	[Ctrl] + [Tab]
Select Location Bar	[Ctrl] + [PageDown]
Complete .com Address 1	[Ctrl] + [L]
Complete .net Address 1	[Alt] + [D]
Complete .org Address 1	[Ctrl] + [Shift] + [Enter]
Open Address in New Tab 1	[Ctrl] + [Enter]
Open File	[Alt] + [Enter]
Select Search Bar	[Ctrl] + [O]
Select Next Search Engine in Search Bar	[Ctrl] + [K]
Select Previous Search Engine in Search Bar	[Ctrl] + [I]
Open List of Search Engines in Search Bar	[Ctrl] + [Up Arrow]
Select Next Entry	[Ctrl] + [Down Arrow]
Select Previous Entry	[Alt+ [Down Arrow]]
Toggle Checkbox	[Alt+ [Up Arrow]]
Select Next Auto-Complete Entry	[Spacebar]
Select Previous Auto-Complete Entry	[Down Arrow]
Delete Individual Auto-Complete Entry	[Up Arrow]
Select All	[Shift] + [Del]
Find in This Page	[Ctrl] + [A]
Find Again	[Ctrl] + [F]
Find Previous	[Ctrl] + [G]
Quick Find Text	[F3]
Quick Find Link	[Shift] + [F3]
	/
	'

3.4 SeaMonkey



A screenshot of the SeaMonkey Browser

If you are a fan of Netscape Navigator, or just want to relive the good ‘ol days, SeaMonkey is just the application for you. The browser is being developed as a community effort to continue from where Netscape Navigator left off. The name ‘SeaMonkey’ was initially used as a pre-beta code name for the Mozilla Application Suite. When the Mozilla foundation announced that they would only focus on the Firefox Web browser and the Thunderbird e-mail client and close all other competing projects, the Mozilla Application Suite community—which was building a browser and e-mail client rolled into one—branched off to form the SeaMonkey Application Suite. The renaming was necessary to differentiate itself from the Firefox and Thunderbird products. The Mozilla community on its part continues to provide the Internet space and resources for the project.

The SeaMonkey suite consists of a Navigator (the browser from the Netscape family), a mail and newsgroups program, an IRC client (ChatZilla) and an HTML editor. It uses the same Gecko rendering engine and similar back-end code that is used by Firefox to

display Web pages. Due to this, it is possible to use Firefox extensions in SeaMonkey too. Like FF, this browser too runs on Windows, Linux or Mac.

When you start the Navigator, browser you can see a few icons on the bottom left edge of the window. These are the links to go to the various applications in the SeaMonkey suite. You see these buttons replicated across all the applications in the suite.

The suite has all the features of the FF browser and even the shortcuts are similar. Since it shares the extensions database with FF, there is a wide variety of plugins to choose from and if you get bored of the Navigator skin, you can install any of the five skins available for the browser. To download any of them, go to <http://tinyurl.com/2h9fdt>.

The similarities however, end there. The SeaMonkey suite has a functionality called Roaming Profiles that allows you to create profiles for individual users. These profiles store the history, passwords, e-mail accounts, and custom settings of each user using the suite. The biggest advantage of this feature is that you can change profiles on the fly without needing to log into a different profile or closing down the browser. Apart from that, SeaMonkey displays messages when moving from secure to non-secure Web pages. The suite is designed to be used by users who normally use a browser and an e-mail client in tandem. The reasoning behind this is that FF and Thunderbird—or any other mail client and browser for that matter—will use more memory than SeaMonkey.

The Preferences settings are located under Edit > Preferences. From here, you can set changes to either Navigator, ChatZilla or any other of its applications. Another strong feature supported natively by SeaMonkey is Roaming Users. This functionality allows you to create a profile which you can use from any install of SeaMonkey in the world. The profile records your browsing history, settings, cookies, address book changes, message views and passwords, among other settings and saves them on Roaming Access Server online when you close the suite. You can install the browser from the Fast Track section of this month's CD.

3.5 Flock

A browser created for the social revolution taking over the Web, Flock allows you to do all your social activities and more, without having to visit the Web sites themselves.

The browser uses the Gecko rendering engine used by Firefox and SeaMonkey. On installation, it allows you to import all the preferences, cookies, passwords and history from IE and Firefox—if installed. The browser supports tabs, RSS feed reading, has a password manager, a pop-up blocker and all the features of the FF browser.

3.5.1 The Flock Toolbar

The uniqueness of the Flock browser lies in the Flock Toolbar which supports all its social capabilities. There are nine buttons on this toolbar which do many things. Let's look at what all of these buttons do:

The first button on this toolbar is the My World button. This button opens up a new tab which displays all your most frequently visited Web pages along with the latest RSS feeds, all in one window. From here, you can quickly access all your most visited sites and read all your updated feeds. This function is like a summary of your surfing habit on the Web.

The People Sidebar button opens up a horizontal column at the left of the window. It allows you to access accounts from FaceBook, Twitter, YouTube, Flickr, del.icio.us and many other Web sites from directly within the browser.

A button called Media Bar opens up a horizontal column across the top half of the browser. This column allows you to view pictures and videos from any page of the internet. This function works best with the People Sidebar.



The Media Bar displaying media updates from FaceBook.

When any of your contacts add images or video to their Facebook,

Flickr or YouTube gallery, you can open them in the Menu Bar even without going to their Web page.

On the same toolbar is the RSS Sidebar button, which allows you to read RSS feeds from within the browser. Whenever the browser detects an RSS feed on a page, it shows you a drop-down window which allows you to subscribe to that feed. You can manage all your subscriptions from the RSS Sidebar and also arrange them in specific folders. The feed preferences button on the RSS Sidebar allows you to choose if new feeds are to be read in Flock or in a different RSS reader like Google Reader.

The fifth button on the Flock toolbar is Favourite Sites. This will open a new column showing you all the sites you have bookmarked with the browser. It also supports a functionality to integrate with <http://del.icio.us> or <http://ma.gnolia.com> and store all your bookmarks there.

Apart from that, another function called Web Clipboard allows you to store snippets of text, images or flash from Web pages you visit. It can store a snapshot of the entire Web page. To store these, you have to highlight the text area, image or Flash object and drag-and-drop it into the Web Clipboard column.

If you are a blogger, there is more good news. Flock supports all the major blogging sites on the Internet, from Blogger, WordPress, LiveJournal to other lesser-known sites. You can configure Flock to remember your blogging account and post directly from the browser. To do so, you have to first configure your account; let's assume that you have a Blogger account and configure the same in two easy steps:

1. Click on Blog Editor on the toolbar and select Setup Account to create an account.
2. Now, click on the Blogger link on the left corner window. This will open up the login page for www.blogger.com. Once you enter the username and the password on the page, your blog settings will automatically get configured with Flock. From then on, you can post entries to your directly from the browser window.

You can use the same method for other blogging services as well. To add a post entry to your blog, you click on Blog Editor from the toolbar and compose your entry in the new window. This window

will allow you to format your text just like Microsoft Word. There is a tab called Source which will allow you to see the HTML source code behind the entry. The Preview tab will allow you to see how the entry will eventually look like on your blog.

The last button on this toolbar is the Photo Uploader. This button too does what it says—uploads photos and videos to major photo hosting sites like Flickr, Photobucket, Piczo and YouTube. You can also add images to your Facebook profile. After you click the button on the toolbar, you will see a new window in front of you into which you can drag and drop images files or videos from your computer. You can then enter a description and add tags to the photos. You can also crop or rotate them from here. Once you have the desired photos and descriptions in place, you can select the account you want to add these photos to and click on the upload button in the window. You can also add the same description for all the photos being uploaded by clicking on the 'Batch' tab above.

3.5.2 Tips And Tricks

One of the biggest advantages of Flock is that it is based on the FF browser. This means that it'll be able to run all the extensions and plugins from the FF add-ons repository. You can also use the FF browser tweaks we mentioned earlier to optimise Flock. Here are the keyboard shortcuts you can use with the browser:



Command Shortcuts

Back	[Alt] + [left arrow]
Caret browsing	[F7]
Change tabs	[Ctrl] + [Tab]
Close tab	[Ctrl] + [W]
Close window	[Ctrl] + [W]
Complete .com address	[Ctrl] + [Return]
Complete .net address	[Shift] + [Return]
Complete .org address	[Ctrl] + [Shift] + [Return]
Copy	[Ctrl] + [C]
Create new collection (in favourites manager)	[Ctrl] + [Shift] + [N]
Cut	[Ctrl] + [X]
Decrease text size	[Ctrl] + [-]
Delete selected auto complete entry	[Shift] + [Del]
Downloads	[Ctrl] + [J]
Find (in page)	[Ctrl] + [F]
Find again	[F3] or [Ctrl] + [G]
Find previous	[Shift] + [F3] or [Ctrl] + [Shift] + [G]
Forward	[Alt] + [right arrow]
Increase text size	[Ctrl] + [+]
New tab	[Ctrl] + [T]
New window	[Ctrl] + [N]
Open accounts and services	[Ctrl] + [Shift] + [A]
Open blog editor	[Ctrl] + [Shift] + [B]
Open blog post	[Ctrl] + [Shift] + [O]
Open favourites	[Ctrl] + [Shift] + [F]
Open file	[Ctrl] + [O]
Open media bar	[Ctrl] + [Shift] + [M]
Open news feeds	[Ctrl] + [Shift] + [N]
Open in new tab	[Option] + [Return]
Open properties dialog (for selected favourite in favourites manager)	[Ctrl] + [I]
Open uploader	[Ctrl] + [Shift] + [U]
Open web clipboard	[Ctrl] + [Shift] + [C]
Page info	[Ctrl] + [I]
Page source	[Ctrl] + [U]
Paste	[Ctrl] + [V]
Print	[Ctrl] + [P]
Redo (after Undo)	[Ctrl] + [Shift] + [Z]
Reload	[F5] or [Ctrl] + [R]
Reload (override cache)	[Ctrl] + [F5] or [Ctrl] + [Shift] + [R]
Restore text size	[Ctrl] + [O]
Save page	[Ctrl] + [S]
Select all	[Ctrl] + [A]
Select location bar	[F6] or [Ctrl] + [L]
Select tab (1 to 9)	[Ctrl] + [1], [Ctrl] + [2], etc.
Star current page	[Ctrl] + [D]
Star and tag	[Ctrl] + [Shift] + [D]
Stop	[Esc]
Topbar menu	[Ctrl] + [Shift] + [T]
Undo	[Ctrl] + [Z]
Web search	[Ctrl] + [K]

3.6 Opera

The Opera browser's tagline, 'Your Web Everywhere', symbolises the fact that this is the only browser in the world that can run on the most number of platforms, from the usual Windows, Linux and Mac OS platforms to mobile phones, the Wii gaming console and the Nintendo DS. Though primarily a desktop browser, Opera has a dominant market share on mobile phones.

Opera Software claims that their browser is "the fastest browser on Earth." This fact is reconfirmed in speed tests run by several independent browser reviewers with the best part of the browser being its JavaScript execution, which performed twice as fast as other browsers.

3.6.2 Features

In spite of the browser not supporting extensions like FF, there are many innovative features built into Opera and extension developers of the FF browser copy quite a few of its features. Opera comes with a variety of in-built features like an e-mail client, an RSS feed reader, a download manager, a BitTorrent client, phishing protection, tabbed browsing, a popup blocker, an ad blocker, mouse gestures, a voice recognition software and the ability to support widgets, which are small programs that run on top of a program or Web page and provide information to the user. One of the newest features of Opera is Speed Dial, which allows you to add up to nine links to a page which is displayed when a new tab is opened. Opera also has a history of being one of the most standards-compliant browsers on the Web.

When you use the browser for the first time, its interface needs some getting used to, but once you have passed that hurdle, you realise that the browser has the most flexible toolbar layout.

3.6.3 Tips And Tricks

There are quite a few settings in the browser that you can tinker with to make it work better.

You can provide nicknames for your most visited Web pages so that you can open them quickly from the address bar. To do so,

click on Bookmarks > Bookmark Page... In the new window that opens, click on Details and enter a nickname under the field.

Now all you have to do to visit this site is to enter the nickname in the address bar and Opera will load the relevant site for you.

You can also remove specific Web pages from your browsing history. To do so click Tools > History. In the new window that opens, enter the Web page you want to delete in the Quick find field and right-click on the results to delete them. To delete your entire history, go to Tools > Delete Private Data.

If you want to visit any URL (which is not already a link) on a Web page, you can highlight it, right click and select Go to URL.

Opera will automatically block most ads for you, however if you want it to block ads that have slipped past its filter, you can right click on it and select Block Content, then select each of the ad images you want to block.

This was also the first browser to support mouse gestures—ability to navigate the Web by performing various actions on the screen using the cursor. You can use these to go back in your browsing history, open or close a tab, go to the next page, save a page etc. Opera also allows you to create your own mouse gestures. To do so, go to Tools > Preferences > Shortcuts and click Edit under Mouse Setup. From here, you can define the shortcuts you want to make.

Like the about:config tool of FF, Opera too has the same functionality. To increase your browsing speed, type opera:config (about:config works too) in the address bar and enter 'connections' in the search field. Now change the value of Max



Adding a nickname to your favourite sites.

Connections Total to 32 and restart Opera. You can also reduce the size of the visited pages History to 100 to speed up your browsing. You can also go to Tools > Preferences > Advanced > Browsing > Loading and change the setting to Redraw instantly. This will tell Opera not to wait for the images or text to load completely before displaying it.

There are tons of other tips available online, which help you modify and enhance the functionality for Opera. A good place to start would be to visit <http://my.opera.com/Tamil/blog>.

3.6.4 Keyboard Shortcuts

Opera has received a lot of dissension from users who complain that their shortcuts are different from Internet Explorer, which makes it cumbersome for most users. Here are some of the shortcuts used in the browser.

Keyboard Shortcuts

Tab	
Open a new tab	[Ctrl] + [T]
Close a tab	[Ctrl] + [W]
Cycle between tabs	[Alt] + [Page Down] or [Ctrl] + [Tab]
Scrolling	
One line up	[Up Arrow]
One line down	[Down Arrow]
One character to the left	[Left Arrow]
One character to the right	[Right Arrow]
One screen up	[Shift] + [Space] or [Page Up]
One screen down	[Space] or [Page Down]
One screen to the left	[Ctrl] + [Page Up]
One screen to the right	[Ctrl] + [Page Down]
Go to top	[Home]
Go to bottom	[End]
Navigation	
Next link up	[Shift] + [Up Arrow]
Next link down	[Shift] + [Down Arrow]
Next link to the left	[Shift] + [Left Arrow]
Next link to the right	[Shift] + [Right Arrow]
Activate link	[Enter]
Activate link in new tab	[Shift] + [Enter]
History	
Previous page in history	[Z] or [Ctrl] + [Left Arrow] or [Alt] + [Left Arrow]
Next page in history	[X] or [Ctrl] + [Right Arrow] or [Alt] + [Right Arrow]



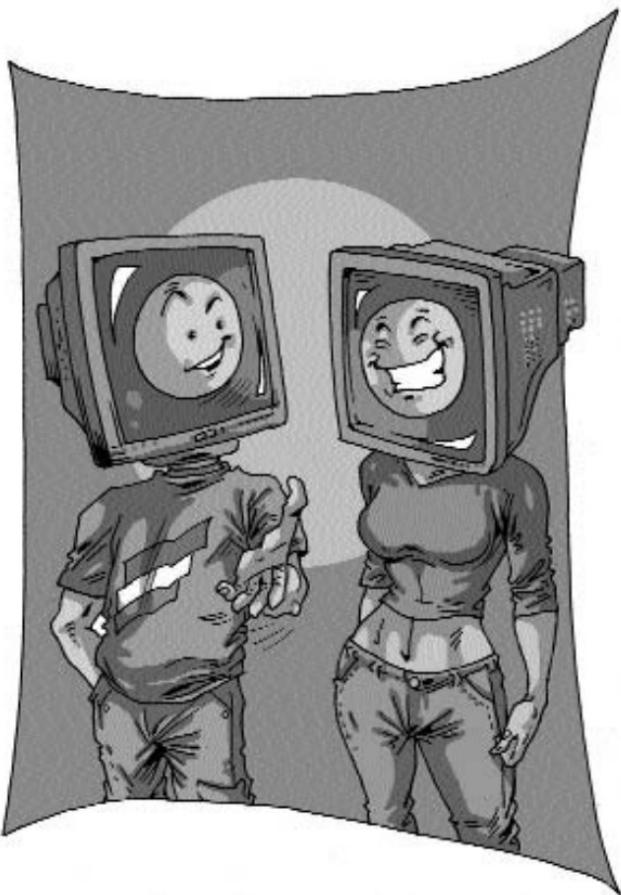
Keyboard Shortcuts

Reload the current page	[F5] or [Ctrl] + [R]
Rewind	[Shift] + [Z]
Fast forward	[Shift] + [X]
Show entire backward history	[Alt] + [Z]
Show entire forward history	[Alt] + [X]
Find	
Find text in page	[Ctrl] + [F] or [,] or [/]
Find next instance of text	[Ctrl] + [G]
Find previous instance of text	[Ctrl] + [Shift] + [G]
Find text in links in page	[,] or [Shift] + [/]
Display	
Use entire screen for viewing page	[F11]
Pretend to be small-screen device	[Shift] + [F11]
Zooming	
Zoom in 10%	[+] or [0]
Zoom out 10%	[‐] or [9]
Zoom in 100%	[Ctrl] + [+] or [8]
Zoom out 100%	[Ctrl] + [-] or [7]
Restore zoom to 100%	[6]
System keys	
Open file	[Ctrl] + [O]
Save copy of page	[Ctrl] + [S]
Print preview (toggle)	[Shift] + [P]
Print page	[Ctrl] + [P]
Close Opera	[Ctrl] + [Q]
E-mail keys	
Basics	
Check e-mail	[Ctrl] + [K]
Write new message	[Ctrl] + [E]
Save message draft (continue later)	[Ctrl] + [S]
Send composed message	[Ctrl] + [Shift] + [S]
Navigation	
Go to next e-mail	[J]
Go to previous e-mail	[U]
Go to next unread e-mail	[H] or [Shift] + [J]
Go to previous unread e-mail	[V] or [Shift] + [U]
Scroll up in mail body (even when focus in mail list)	[Alt] + [Up]
Scroll down in mail body (even when focus in mail list)	[Alt] + [Down]
Toggle view: message list/message body/list and body	[I]
Threads	
Expand current thread	[Right Arrow]
Close current thread	[Left Arrow]
Expand all threads	[Shift] + [Right Arrow]
Close all threads	[Shift] + [Left Arrow]
Marking	
Mark selected e-mail as read	[K]
Mark selected e-mail as read, and go to next unread e-mail	[G]
Mark selected e-mail as read, and go to previous unread e-mail	[T]
Mark selected e-mail as unread	[Shift] + [K]
Mark all e-mail in active view as read	[Ctrl] + [Shift] + [A]
Responding	
Reply to message	[R]
Reply to all recipients of message, including sender	[Shift] + [R]
Forward e-mail	[F]
Redirect e-mail	[D]

There are a lot other browsers out there. Most of them perform the same functionality as those mentioned above. While Internet Explorer will offer you the best Web site rendering experience, it is up to you to choose between speed and Web site compatibility. If you can forego the latter, Opera should be your best bet.

The browser to look out for in the future is definitely Flock. With a robust set of features and the idea of developing relationships based on your surfing habits makes this browser one to keep an eye on.

IM and VoIP



Instant Messaging and Internet Telephony services work, unlike e-mail, in real-time. IM is free, as are voice calls from computer to computer, while calls from computer to phone—local or international—are very cheap. IM services typically use a client on the local machine to connect people, though some services like Gmail and Yahoo! Mail offer the option of online chat without having to install software. Here is a walkthrough of popular chat clients and telephony services; just don't be at it while in office—chatting, especially, is not popular with most bosses!

4.1 Chat Clients

4.1.1 Google Talk

Among the major players in this segment, Google is a comparatively late entrant. It started off with a downloadable client called Gtalk, which enabled users to chat and make voice calls. Google also enabled the chat functionality from inside Gmail. This proved to be wildly popular because people did not need to do messy installs and configure settings: just click on the contact once you were in, and start talking with your Gmail contacts.

After downloading the client from www.google.com/talk, select Settings > Connection and enter the network settings. Sign in using your Gmail username and password and you are all set! Google uses the open source Jabber protocol for its chat services.

For configuring the look and feel of the chat window, select Settings > Appearance and choose the chat themes from the drop-down. Each theme has a preview; choose the one you like and click OK. Another way to change the way your contacts are displayed in your Friends list is by using the View button. You can enable all your contacts to be seen even when they are offline, or to only see online contacts. Additionally, you can see all your online contacts or just a few of them, by clicking on the Show one page option.

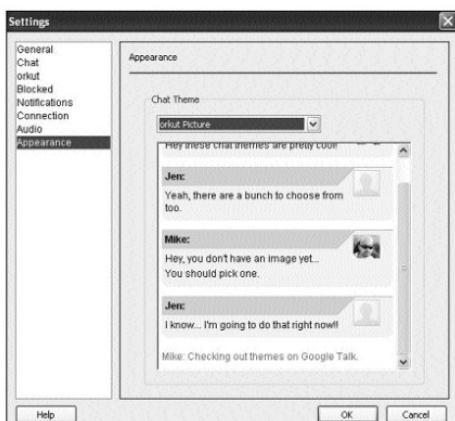
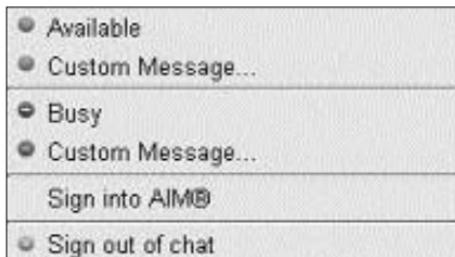
Gtalk is tightly integrated with Orkut. If your contact is online on Orkut, you can scrap him / her from inside Gtalk. Viewing the Orkut profile from inside Gtalk is also possible. If the contact is not online, you can send voicemail messages. These will appear in the Gmail Inbox of the receiver, and can be heard as an MP3 file. Any voicemail pending for you is indicated by a green phone icon.

For making calls from inside Gtalk, click on any name from your list. If that person is online using Gtalk (contacts can show up in the list if they are logged in to Gmail and has chat enabled), and if the Talk option appears, click on the button and make a call. File transfers in Gtalk are pretty fast, and there is no restriction on file size. The transfer is peer-to-peer and pretty fast.

For chat inside Gmail, there are several features that are not available in Gtalk. One is the ability to hold a group chat (with more than one person); another is the new Flash smileys that comple-

ment Gmail's traditional no-frills ones. Another new feature is the ability to talk with your AIM friends. Click on the inverted triangle beside your name in the chat window and select the "Sign into AIM" to do this.

Gtalk fans have the option of using a Google Talk gadget that can be embedded in Web sites. This widget has tabbed features: in both chat-inside-Gmail as well as in Gtalk, you need to open multiple windows for each conversation. The Talk widget allows you to use tabs such that you don't have to navigate around your screen to click on individual conversations. Just hit [Ctrl] + [Tab] to cycle through conversations.



Different themes that can be used in Gtalk

4.1.2 Windows Live Messenger

Windows Live Messenger is the new-age version of the old MSN Messenger, complete with Vista look and feel. All Hotmail, MSN, and Windows Live accounts can be used for logging in. The latest version is 8, which is available for free download at <http://get.live.com/messenger/config>. The installation takes quite a while; you would have to download some critical updates from Microsoft Update before you can start chatting.

WLM (Windows Live Messenger) has refined the basic chat features of its earlier avatar while adding some new ones. The messenger window has a number of vertical tabs, with each tab letting you use a different feature—including News, Dating, Games, and

Video. Clicking on some tabs will, however, open new windows in your browser. The number of contacts has been increased from 300 to 600, while a word wheel enables you to quickly search the list. This tool looks not only at names, which can keep changing, but also looks at other data like personal status messages, phone numbers, nicknames, and e-mail IDs. These details can be checked by hovering the cursor over each contact. Notes can be added for each contact, while clicking on the name automatically allows you to start writing mails.

Then there are other features like IM offline contacts, or changing your status to remain hidden, or time-stamping each message. (Use Tools > Options > Messages > Show Time to turn this on.)

WLM has file-sharing capabilities as well. Just drag a file into the message window and it will get shared with the contact. A complete redesign has been effected, with new icons, and the default colour has been set to orange (some other available colours are red, pink, blue, and yellow). The colours can be changed by clicking on the paint-brush icon at the top of the messenger window. Apart from smileys, winks can be sent to your contacts during a conversation; get them from Tools > Winks. If you're bored of the same bland chat window, you can select a custom background using pictures stored on the local machine. Get this at Tools > Backgrounds.

Windows Live Messenger can be used to make voice calls to other PCs or phones to anywhere in the world. Called Windows Live Call, PC-to-PC calls are free to your contacts. Video calls are also supported in Windows Live Messenger, and even if you don't have a webcam yourself, you can always see



You can choose custom backgrounds for chat windows



You can use plenty of services in Live Messenger

your friends talking to you in a small window. Search is integrated with the phone service, which makes it easy to find numbers.

The extra services that come with Windows Live Messenger are perhaps the most attractive (or most annoying, depending on point of view) feature. Just by opening your chat window you can bid on eBay, visit Xbox Live (and also see your contacts' status if they are on Xbox Live), get news, chat with bots, play games with friends, listen to music, invest in stocks, find a date at Match.com, watch free Live videos...

When it comes to interoperability, Messenger users can natively chat with Yahoo! Messenger users, while they can be accessed by most major multi-protocol clients.

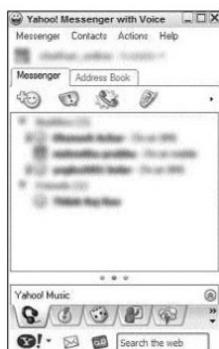
Apart from sharing files, you can also collaborate on a document. Changes you make to a document are immediately updated and viewable. Anti-virus protection is incorporated with the service, and files are encrypted before being sent. There is no limit on file size, and if the power goes off, the next time you log on, the transfer will resume from where it had stopped.

4.1.3 Yahoo! Messenger

Yahoo! Messenger works best on Windows, with the latest version being Yahoo! Messenger 9 Beta for Windows. For Linux, rpm and deb packages are available for older versions, while for Mac, the latest stable version is Yahoo! Messenger 2.5.3. Here we'll discuss the Yahoo! Messenger features for Windows.

Yahoo! Messenger can be installed in two ways: either you download the client from www.messenger.yahoo.com/download.php and install it, or use a Web-based installation from <http://webmessenger.yahoo.com>. For the second option you need to be online, as most of the data needed for installation is based on the server. This will take some time, so we suggest you go with the offline installation.

The first thing to do after installation is to key in the network settings. Select Messenger > Connection Preferences and click on



All your contacts at one place

the appropriate radio button. Once that is done and you are online, you can start chatting. Yahoo! Messenger has a slew of features, which sometimes seems a bit of overkill. You have a whole load of animated smileys for

every mood under the sun, and then some. There are specific themes called IMVironments; you can set these while chatting with friends. Unlike Gtalk's chat themes, both you as well as your friend will be able to see the selected theme, so don't be surprised if you suddenly see the chat window swimming with fishes—it's the fish-tank theme. Some IMVs are interactive, allowing both of you to draw doodles using the cursor.

Another fun feature is the ability to send sound bytes in a variety of languages and in different moods. The most fun are the Insults, where you can send quite a few insults, some in languages other than English. Using a feature called Soundboard, you can also insert sounds like bumps and groans to insert that extra realism into chats. The Soundtrack feature will allow you to let a choice MP3 file play while you are chatting, and your friend will also hear the track if he or she wishes to.

Apart from enabling chats, you can also make voice calls from PC to PC or PC to phone. Available only for Windows, these VoIP features are integrated under Yahoo! Voice. You can access Yahoo! Voice from the main chat window after logging in. PC to phone calls need credit (a call to the US costs \$0.01, or about 40 paise per minute), which comes in packages upwards from \$10 (Rs 400). Get the complete list from http://voice.yahoo.com/phone_rates.php. Yahoo! Voice has a Phone In and Phone Out: Phone Out is the regular voice calls you can make from your PC to any phone, mobile or landline in over 30 countries at very low rates. Check out the rates at http://voice.yahoo.com/phone_calls.php.

Phone In is a feature where your computer gets a normal phone number, which can be called from any other phone. You just accept the incoming call and start talking. If you're tired of



Themes (called IMVironments) in Yahoo! Messenger

making the effort of accepting calls every time, you can use OpenTalk, which enables your PC to accept a call from another PC without you clicking anything.

It is also possible to share files with other users in Yahoo! Messenger. This feature has been enabled from version 8 onwards. Firewalls must be configured so that the transfer is not blocked. Sending is as easy as dragging the file into the IM window. There is no limit on file size. You can share up to 300 photos with each contact. Click on the Photos button at the top of the IM window to share photos with them. You can view the shared photos in a box that expands near the IM window.

If the firewall settings allow, webcams can be used in Yahoo! Messenger. You can also enter public chatrooms, or create your own chatroom and invite friends.

Yahoo! Messenger has loads of plugins available for various purposes. One is Launchcast Radio, which lets you listen to songs in many genres from pop to rock to country to Indian from hundreds of radio stations around the world. Other plugins let you play games, access news and weather information, and get updates from your contacts' blogs. These plugins need to be installed; click on the Add Plug-in button in the Messenger window to install a plugin. Some plugins let you customise the experience when you chat, like the ability to check your calendar and get news feeds.

Yahoo! Messenger interacts with Windows Live Messenger, albeit in a very basic barebones way. You can only do plain chat; features like voice calls, sharing photos, IMVironments, etc. are not supported.

4.2 Multi-protocol IM Clients

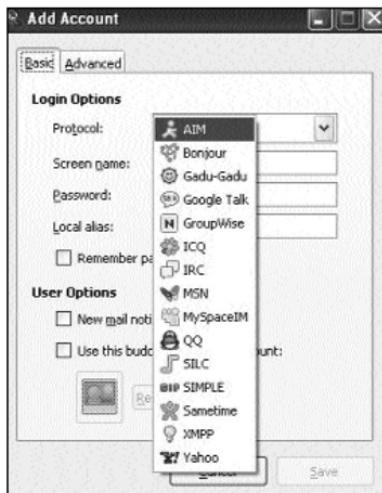
This category of software is a boon to people who have IM accounts with different services like Google, Yahoo, and Microsoft. These clients enable you to log in to multiple IM services from the same place. There are both online clients as well as offline clients. Online clients include Meebo; we will only talk about locally-installed software.

4.2.1 Pidgin

Pidgin (formerly Gaim) is a cross-platform, free, and open source chat client that works across Windows and Linux. Download it from www.pidgin.im/download. Pidgin works with a wide variety of services, right from the major ones like AIM, Yahoo, Google Talk, MySpaceIM, and MSN to the more exotic ones like ICQ, IRC, Gadu-Gadu, Bonjour, QQ and Zephyr. The latest version is 2.3.1. Pidgin can also be used in a non-graphical console mode; this mode is called Finch.

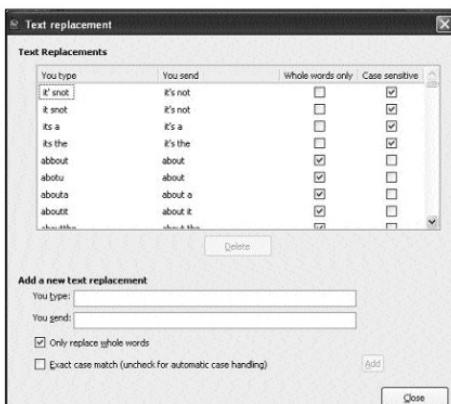
Pidgin integrates well with window managers in operating systems, enabling it to be operated from the System Tray. Once installation is done, right-click on the Pidgin icon and click on Accounts. After clicking on Add, choose the service you want (Gmail users would choose XMPP here) and enter your login and password. You can add as many accounts as you want, because Pidgin is sure to support almost all your IM accounts. When you fire up Pidgin, you will be prompted to enter all your passwords for the individual services. All your friends who are online will be displayed in your Buddy List, arranged either alphabetically or categorised as services.

Pidgin incorporates quite a slew of features; it supports most standard features like emoticons, notifications for different actions, file transfers, and tabbed conversations in a single window. However, some features go beyond what is offered in other similar types of software: you can customise sound alerts for different actions like a friend coming online or going offline. Go to Preferences > Sounds and check out the options. Conversations can be stored in a log file, and the misspelt words in a conversation can be highlighted.



Pidgin supports more protocols than you'll ever need

What really is the strong point for Pidgin is plugins. Right-click on the Pidgin icon on the taskbar and click on Plugins. Read the descriptions right there and select the plugins you want. Some add basic functionalities like the ability to store notes about your contacts, or configure how long after you should be shown idle. Some are smart: the Text Replacement window keeps checking for user-defined misspellings and sends out the correct phrase, also defined by the user. So for example, if you have this habit of typing “it snot” instead of “it’s not” and are fed up of getting emoticons from friends wondering about stuff better not thought out, this is a lifesaver. Other supported features are encrypting chats using RSA and grouping multiple contacts under a single name so that group messaging becomes easy.



Text replacement rules, for those key misses

4.2.2 Trillian

This is a proprietary multi-protocol IM client that comes in two versions: Basic (free) and Pro (for a fee—\$25, or Rs 1,000). It is currently Windows-only (XP and Vista). It supports multiple chat services like AIM, ICQ, Windows Live Messenger, and Yahoo! Messenger for both Basic and Pro versions, while additional support for IRC, Novell GroupWare Messenger, IRC, Jabber, and Skype is available in the Pro version. Trillian is named after a character of the same name in Douglas Adams’ *Hitchhiker’s Guide to the Galaxy*. The current version is 3, while a new version called Trillian Astra is in alpha testing stage. This version can be used from inside the browser, which will make it work in other operating systems. Download the free version from www.trillian-messenger.net.

Trillian has functionalities from basic chat to more advanced voice chat and webcam support. Apart from supporting multiple

connections to different chat services, it manages multiple accounts of the same chat service. Trillian represents different chat services using different coloured spheres: blue for Windows Live Messenger, red for Yahoo! Messenger, purple for Gtalk (Jabber), etc. It enables users to access services like file sharing and chat room joining (if they are available on the parent network). Configuring the network settings for these services, especially when you want to use webcam and voice support, is a bit dicey and varies depending on the network.

Each service is run by a specific plugin whose settings can be changed in terms of ports. For example, for file transfers in MSN, use port 6891, which can be set from MSN plug-in Preferences > Miscellaneous > Ports used for outbound file transfers. (This feature is only available in the Pro version.)

Trillian's look and feel is like Windows in terms of skins and navigation between preferences.

Additional skins can be downloaded from the official Web site (www.ceruleanstudios.com/downloads/downloads.php?t=s) as well as from deviantART (www.deviantart.com) and Customize.org (<http://customize.org/trillian>). Trillian uses its own skinning engine called SkinXML. You will get a wide variety of brilliant skins from these resources.

Metacontact is a very useful feature in Trillian's contact management; it prevents duplication of contact details. Available



Transfer files securely in Trillian



Managing your contacts in Trillian

in the Pro version, Metacontact lets you collect the different account names of the same person in different networks under one entry. The names are in tree format, which can be expanded or collapsed. Saving records of conversations or other activities is strongly supported in Trillian: log files can be stored either in text format or in XML. Saving history in XML is particularly useful as chats can be easily searched.

Instant lookup is a great feature in Trillian which ensures that you do not need to display your ignorance any longer. Say you're talking with a friend of yours who is a bit of a history buff. So when in the middle of discussion about a recently-released movie he says Operation Barbarossa was one of the greatest mistakes in history, you need not scratch your head: Trillian uses real-time referencing against Wikipedia, and such names and terms appear with a green underline and are looked up in Wikipedia for a brief definition. So even if you didn't know what Operation Barbarossa was, you could chip in saying Hitler was probably acting on astrologers instead of his army commanders...

Emotiblips is another unique feature of Trillian, which can be described as the video equivalent of an emoticon. Emotiblips lets users send songs or video clips to contacts during video chat sessions. Supported file formats are MP3, WAV, WMV and MPEG. File transfer is available in Trillian, too, and you can configure your anti-virus software to scan all incoming files.

4.2.3 Miranda

A light and customisable IM client designed for Windows, Miranda is available for free under the GPL. Miranda is basically a graphical interface and plugin architecture where all features are installed as plugins. This enables the software to consume minimal system resources and yet provide desired functionality. It supports a large number of IM protocols services like Yahoo!, Jabber, MSN, AIM, Gadu-Gadu, Skype, etc. Get Miranda from <http://www.miranda-im.org/download/>. The downloaded package provides support for major protocols, while addition protocol support can be obtained by downloading add-ons from <http://addons.miranda-im.org/>. Whatever service you are using, you

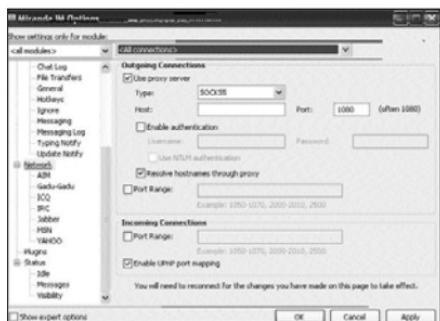
will be sure to get the plugin for it. The latest stable version is 0.7.3.

Miranda is chock-full of options, which is sure to intimidate a casual user. After installation, right-click on the Miranda icon in the taskbar and select Main Menu > Options. Click on the Network item on the tree and key in your connection settings. For configuring the settings for individual services, you'll need to double-click on each item in the tree view and then enter individual details. Each network will have its own, unique options, and you may have to download additional plugins to enable some functions: for example, for logging on to Gtalk using Jabber, you would need to install an SSL add-on. Navigating through these items and configuring them to your requirements could take up time, but it will be worth the while.

Out of the box, the Miranda install comes with a fraction of functionality, which can be expanded by installing plugins, but even this functionality is much more than what is available in other clients. First up are options for customising appearance. Different types of icons and status messages are available for each and every protocol present or installed. The menu items in the chat window can also be customised to give either an austere, minimalist look or a much more detailed one. Similar is the case with sound alerts: you can have sound alerts for every event that can possibly take place in an online chat, or you can shut all sounds out. If you are a chatroom regular and kind of dissatisfied by the limited configuration options in your client, then prepare to be blown away by the rash of options under Events > General. A list of options avail-



Accessing the Options menu item



Configuring network parameters in Miranda

able to the click-happy user includes using timestamps, sounds played when users are kicked out of chat rooms, adding a colon after usernames, stripping colours from messages in log files, docking the position of the chat window, different icons for user actions, and...

Plugins for Miranda are usually in the ZIP format, and need to be unzipped in the Plugin folder inside Miranda. After the plugin has been installed, open Miranda's options and navigate to Plugins. If the newly-installed plugin is present in the list, everything is set. Some things to watch out for is to always use the latest version of the plugin, and not using a Unicode plugin with an ANSI Miranda build. Most plugins come with their own readme files, where a detailed description on how to install the software would be provided.

There is a large list of available plugins at <http://addons.miranda-im.org/>. They are divided into broad categories like protocols, history, security and privacy, database, messaging, interface, etc. There's IM-History Client Suite, which allows you to store your chat history online securely; RSS News, which lets you read RSS feeds in Miranda; Chess4Net, which lets you play chess—these are just three of the thousands of plugins available!

4.3 VoIP

We have featured two VoIP software here: the fairly well-known Skype, and the comparatively new Gizmo. These services use peer-to-peer networks to make calls to phones and computers.

4.3.1 Skype

Skype runs on a variety of platforms: Windows, Linux, Mac OS X, Windows Mobile, and also phones like the Nokia N800/N810. The software is free for download (www.skype.com/intl/en/).

Skype calls can also be made from special phones, so you are not tied to your computer for making cheap long distance calls. The latest versions of Skype are 3.6.0.244 for Windows, 2.6.0.184 for Mac OS X, 1.4.0.118 for Linux, and 2.2.0.36 for Windows Mobile.

Once you have downloaded and installed the software, you will be prompted to log in, or create a new account. You just need to fill up a small form by choosing your account name, password, e-mail ID, and country. Simple, fast, and hassle free; of course, you also need to enter the network settings under Tools > Options before you can connect. After you connect, you will be prompted to make a test call to test your headset and microphone. If you hear the automated voice and make a test call, your equipment is working fine. Otherwise you might need to install newer hardware, or install the latest patches for your older hardware.

Once these preliminaries are done, click on the Contacts tab and use the Add button to add contacts. Enter the Skype name of your contact and click on Search. You can also search by e-mail ID or name, or add a phone number for making calls to ordinary phones. You can also send SMSes to other phones. This phone number can be anything—local, national, international; landline or mobile. The calls are not free, though they are very cheap, especially international calls. For example, calls to the US cost around Rs 1.15 per minute. Credits can be purchased in packages of 10 (Rs 570) using either PayPal or your credit card. This feature is called SkypeOut.

Another feature is SkypeIn, where you can make a call to a computer from your phone. SkypeIn is available for numbers in countries including Australia, the UK, the US, and France. The SkypeIn system works this way: suppose a user from the UK wants to talk to a Skype subscriber in the US. That subscriber will take a UK telephone number as a SkypeIn number, and the UK caller will need to pay local rates for calls made to the US from ordinary phones. SkypeIn is not available for Indian numbers.

Videoconferencing works very well in Skype. It is currently available for Windows, Mac, and Linux. For this, you need a Skype-supported webcam. In the latest version (on Windows), Skype supports high-quality video with features like full-screen viewing and screen-



Adding contacts in Skype

in-screen. Video calls are currently one-to-one, though you can have audio conferences with up to 24 people simultaneously.

Videoconferencing is free in Skype. Like with other chat clients, you can text-chat with your contacts in Skype, too. Other features like call forwarding and voicemail are also supported.

Skype has a service called Skype Prime Beta, where you can earn money when people call you for help. This feature is very useful for professionals who would like to earn by offering their services. Enter the details in the form provided and you could soon have clients calling up from every corner of the world! Another business-oriented service is SkypeFind, which lets you search for local businesses by name and postal code.

Skype has a series of video broadcasts called Skypescasts under its Live service. These are the voice equivalent of chatrooms where people gather to speak on any common topic. A complete list is available at <https://skypescasts.skype.com/skypescasts/home>.

4.3.2 Gizmo

Gizmo is another peer-to-peer VoIP service. It uses open protocols like SIP (Session Initiation Protocol) and Jabber, unlike Skype, which uses proprietary systems. It works across the Windows, Mac, Linux, and Symbian operating systems. It is free, and can be downloaded from www.gizmoproject.com. The latest version is 3.1.2.

After installing Gizmo, you are prompted to test your headphones and microphone. Go to **Edit > Options** and configure the network settings, add contacts, and configure sounds. For making calls, use the Contact tab in the main window and click on the green button. You will hear a ringing sound, which will indicate that your call is being transmitted.

Gizmo bills itself as a “phone for your computer.” Gizmo, like Skype, can be used for making calls to other computers and



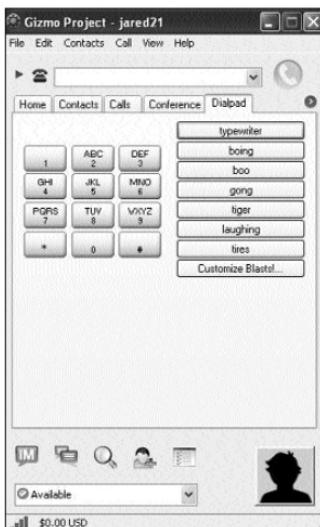
Configuring sound options in Skype

phones irrespective of location. But Gizmo is comparatively cheaper than Skype. The Call Out feature is where you call other phones from your machine, while the complementary Call In feature allows other users to call your computer from their Gizmo phones. (Call In numbers are not available for Indian numbers.) Area 775 number is another feature in Gizmo, where you can get a free Nevada phone number which can be used on a computer. Any calls made to this number can be answered from your computer, with features like caller ID, dual ringing, and customised greetings when you are not available. Group chats are available, as well as free conference between computers, landlines and mobile phones by using "FreeConferenceCalls" and also IM.

Each Gizmo account has its unique SIP number, which lets users from other VoIP networks call Gizmo for free. Apart from the call features, file transfers can also be enabled. For mobile phones that can connect to the Internet, Gizmo can be used for making VoIP calls.

Gizmo5 (www.gizmo5.com/pc) is a free phone software for making cheap international calls. Calls to the US, for example, cost only 75 paise per minute. It comes in two versions, and a full list of handset models is available at the Web site. Credit can be bought in different denominations, and file transfers can also be made, though you need the Full version, not the Lite version. With Gizmo5 Full version you can also IM your AIM, Gizmo, MSN, and Yahoo! contacts.

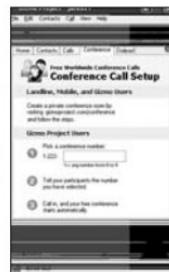
Gizmo has a slew of features not available in Skype. As a chat client, it can interact with other networks that use SIP, some of which are Windows Live Messenger, AIM, Yahoo! Messenger, and Google Talk. It allows call recording and free voicemail. Bloggers who use LiveJournal can make voice posts using Gizmo. The calls



The keypad, for making calls

made from Gizmo-to-Gizmo are encrypted using SRTP (Secure Realtime Transport Protocol), a robust security protocol.

Gizmo users can make free calls to landline and mobile numbers from around the world for free by inviting friends. Each invite will earn you 20 minutes free talk-time. The full list of countries to which calls can be made is available at <http://www.gizmoproject.com/learnmore-allcallsfree.html>. Indian numbers are not included in this list, however. If you use Asterix PBX or another such voice service, you can log in from Gizmo by creating a secondary login. This will simultaneously let you access office calls even if you are not physically present in the office: you only need to be online. Gizmo, however, has no video functionality.



Make a conference call

Security



Fact: as of now, if you install Windows on a new computer and just get online—with neither an anti-virus nor a firewall installed—your computer will most likely be infected by a virus within minutes.

That should get you into the spirit of reading this chapter!

5.1 Introduction

“Security” used to be a synonymous with anti-virus software. Till just a little over a decade ago, most data transfers involved a floppy disk, and activating a virus scan every time a floppy was accessed ensured that a user’s PC remained virus-free. Then the Net exploded.

We don’t need to tell you what the Net is like today: “security” covers a lot more than just viruses. The threats facing a lay Net user include rootkits, browser hijackers, phishers, worms, Trojans, keyloggers, data sniffers, and more. And for a system to be termed secure, all these bases need to be covered.

Unlike the anti-virus packages of yesteryear, it is nowadays common to come across Security Suites—bundles consisting of an anti-virus, a firewall, and anti-spyware in one package.

5.2 Terminology

Here we talk about some essential terms associated with security.

5.2.1. Phishing

This refers to the practice of making a user part with sensitive financial and other personal details by posing as an authorised agent. Usually, this takes the form of an e-mail asking the user to update details at a bank or company Web page, the link to which is provided. The Web page looks just like the original, and a gullible user might enter the requested details, like username and password, to log into his account. These details are routed to the scammer who then misuses it at the real site, sometimes at the same time.

5.2.2. Rootkits

This type of application burst into the public consciousness after it was found that Sony had been using this technique to hide its anti-piracy applications from showing up. Rootkits are applications that can hide other applications from showing up in any of the system’s monitors. By inserting rootkits, it is possible for malware to avoid detection even by anti-virus applications.

5.2.3. Browser Hijackers

These are mostly of nuisance value. Browsers that have been hijacked always open a certain Web page by default on start, or change the home page of the browser, or do not allow certain sites to be visited—usually anti-virus or anti-spyware sites. Trying to uninstall or delete the responsible file proves futile, since some part of the malware is loaded into memory when the system boots.

5.2.4. Spyware

This refers to a class of applications that gains entry into a system and monitors the user without his knowledge. Spyware report back to a remote server and can result in the leaking out of confidential details like usernames and passwords.

5.2.5. Keylogging

This refers to the use of software to record the sequence of keys pressed by the user. Usually, a report of the keystrokes are passed on to the spyware controller on the Internet. Keystrokes of special interest are passwords and usernames, or online banking details.

5.2.6 Viruses

Applications that tag along with other files; running the file causes the virus in it to infect the system.

5.2.7. Trojans

These are applications that gain entry into the user's system by posing as something else, often as free screensavers. After gaining entry into the system and being installed or run by the user, they can manipulate user data. Trojans have different behaviour once they gain access into the system. For example, some allow unauthorised access to a remote user into the user PC by opening up ports, some can act as keyloggers, and some can prevent applications from running—especially anti-virus applications.

5.2.8. Worms

These are applications that autonomously transfer themselves through a network. They do not need an associated file to spread

themselves. Worms can cause harm just by overloading any network by moving from one node to another in huge numbers, thus causing legitimate traffic to be delayed or disrupted—though they can also carry malicious code.

5.2.9. Data Sniffers / Packet Sniffers

These are applications that try to detect the contents of a message you've sent. Any data that is transferred—e-mail, pictures, files—are split into data packets and then sent over the Internet. Any server through which the data packets pass can install a sniffer to detect the contents of the packet. Files that contain binary data like pictures or video can also be deciphered, but e-mail messages and other textual content show up as text in the sniffing application, making them easier to decipher. The encrypting of messages is the only way to prevent this from happening.

5.3 Steps To Complete Security

Absolutely complete security is impossible, but in general, to overcome the various security threats we've mentioned, the following steps need to be taken:

- Unauthorised access to the system can be blocked by firewalls. Both inbound and outbound traffic can be monitored, and any unauthorised access can be checked.
- Viruses and similar programs can be checked with anti-virus applications.
- Spyware, browser hijacking, and more can be blocked by anti-spyware applications.
- Track / History cleaners constitute a special set of applications that specialise in cleaning up a user's browsing history. Windows stores the list of sites visited by a user in many places, including the Registry. It is easy for someone to know where the user has been by looking in those places. History cleaners can make this task difficult.
- Miscellaneous security applications including specialised tools like rootkit detectors and Registry cleaners can be used to augment the capabilities of the larger applications.
- Finally, since most break-ins can be prevented by keeping a watchful eye, the user plays an important role in ensuring that scammers do not have it easy. When accessing public PCs it is imperative that the browsing and other histories are deleted once the session is over. "Social Engineering"—a euphemism referring to, among other things, the use of informal conversations to extract confidential information—plays an important role in breaking usernames and passwords. Users need to be aware of such situations as well.

5.4 Anti-virus Software

5.4.1 Introduction

All Anti-virus (AV) applications operate on similar principles. All have a virus database that contains a list of file patterns that can be used to identify a virus. A pattern refers to a sequence of bits that is unique to a particular file. The second component of an AV is a scan engine which scans files and compares it with the pattern file to verify if there is an infection. While the scan engine can be run when the user explicitly activates it, it can also be set to run automatically any time a file is opened or a program is executed. This latter mode of operation is called a real-time scan.

With an increase in threats from online data transfers, all AV modules have begun incorporating additional capabilities to the basic program that only scans local files. It is rare to find an AV program that does not include capabilities like e-mail scanning (inbound and outbound), downloaded file scanning, spyware scanning, and such.

5.4.2. Avast!

The Avast! Home Edition anti-virus is free for non-commercial use, given that the user has acquired a freely-available registration key from the site within 60 days of installation to continue using the product. Besides the basic virus scanner to protect local folders, Avast! also includes modules to scan online file transfers. There are modules that scan files transferred during regular browsing, peer-to-peer activities, Instant Messaging, communicating through newsgroups, and e-mail. Avast! creates a database of the files in the system when the system is idle. This “Virus Recovery Database” can be used to recover files damaged by a virus.

The free Home Edition only allows using the simple interface



The Avast! Simple User Interface

of the main program window. This interface has a single function—configuring a scan. Here the scope of the scan can be set, along with the type. Archives like ZIP or RAR files can be excluded from the scan if needed. One notable absence in Avast's arsenal is the scan scheduler, which one should consider a basic requirement.

The program's options are available from the context menu of the System Tray icon. Under the Program Settings link, different aspects of the program's behaviour can be configured. Avast! supports different modes of informing the user about a virus detection—through SMTP or MAPI e-mail, through ICQ or MSN Messenger Instant Message, or via a direct alert to a printer. These settings can be made under the Alert link.

The On-Access Protection link allows controlling the behaviour of the many real-time scanning modules currently active.

Clicking on the Details button at the bottom allows each module to be individually configured. Under the Instant Messaging link, the level of security desired can be set, or a custom level can be set by clicking on the Customize button, under which the different networks that need to be monitored can be set.

Similarly, under the Internet Mail link, the level of security can be set or customised. Under the customisation options, rules to



Avast!'s Resident Protection controls



Avast! has options for configuring even what it should do during IM sessions

scan inbound and outbound traffic on SMTP, IMAP, and NNTP (newsgroups) can be set. A heuristic tab allows for checking for suspicious mail by their features rather than any viral pattern. Under this tab, the program can be set to alert the user when any program sends frequent e-mails or mass e-mails, or when the filenames have too many blank spaces—a common technique to hide the actual file extension from showing up (such as “readme.txt.exe”).

Similarly, options for the other modules—P2P, Outlook, and Network Shield can also be set. It needs to be noted that while the Network Shield operates similar to a firewall, by Avast’s own admission, it cannot be seen as a replacement for a dedicated firewall.

The Web Shield module operates as a proxy server, and in Windows XP requires no tweaking of the browser since it monitors port 80, the default port used for Web traffic. In older OSes, browsers need to be configured to connect through it (IP 127.0.0.1, port 12080). This module acts as a preventive measure, detecting malicious files before they are downloaded. A URL-blocking function can also be used to prevent specific sites from sending data.

Avast! offers updates on a daily basis, and most updates are a few KB in size.

5.4.3 AVG AntiVirus

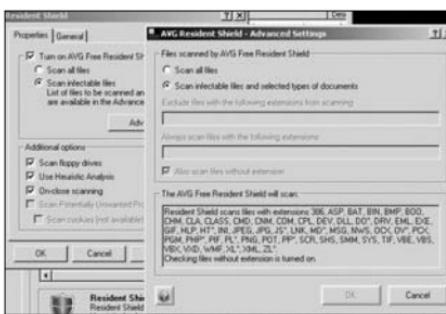
The free version of AVG offers an e-mail scanner along with the anti-virus. There are fewer modules to configure in comparison with Avast!. It includes a Scheduler, which allows a scan to be scheduled at a particular time.

Under the Tests menu, various scan parameters can be specified. By default, AVG automatically heals infected files, and



The AVG Extended Mode control panel

only if a file is incurable does the user have the option to either quarantine or delete the file. The user can choose whether to heal a file automatically during a scan (from the Program Settings, under the Service menu). If this is disabled, AVG does not allow the user to heal any infected file during a scan—the only options are to delete and quarantine. The Heal button in the scan window is greyed out, and may be available only in the paid version. Besides the scanner, AVG also has an e-mail virus-checking module.



The various settings for the Resident Shield of AVG

AVG, too, has updates almost daily, and the process can be managed within the application. The user has two update options: either online updating, or from a folder. This ability to update from a folder is quite useful for people who do not have Net access.

The AVG control panel lists the program components. Under the Resident Shield link, the real-time scanning behaviour can be set. File extensions that can be skipped can be indicated here. This is an important time-saving option, since viruses like virtob only infect EXE files, and the scanning of the rest of the files can be avoided.

5.5 Firewalls

5.5.1 Introduction

Firewalls are devices—software and hardware—that prevent unauthorised entry into a system. Any network interface has 65,535 possible points of connection (these are called ports). When a PC is connected to the network, these ports are accessible by any other PC on the network. This presents an opportunity for those with “malafide” intentions to access the contents of a PC without permission. A firewall usually operates by blocking or monitoring all

ports of the interface. Traffic can be inbound or outbound. When data is sent to a user's PC, it is inbound traffic, and when the user's PC is sending data, it is outbound traffic. Usually, inbound traffic is considered more dangerous since it could be a virus or Trojan trying to gain access into the system. But outbound traffic can also be harmful, since it could be an application that is surreptitiously passing data to someone else—for example, a keylogger could be operational.

5.5.2 Comodo Personal Firewall

CPF has won rave reviews that pretty much say it's better than most other personal firewalls, free or paid.

After installation, a Wizard helps you configure CPF. While the automatic configuration works well, you can also manually change parameters—like designating which network device is used to connect to the Net, which applications are considered trusted, and the frequency with



The Summary window in Comodo Personal Firewall

which CPF should alert you. CPF maintains a database of trusted applications that should be allowed access to the Net; you can either prefer to go with the database, or scan your system for applications and mark them as trusted or otherwise.

Depending on the Alert frequency set during the configuration, CPF will pop up alerts whenever an application tries to access the Net, asking you whether it should allow or deny access. Besides the Application, Component, and Network monitors, which watch installed programs, OS components, and the network interface respectively, CPF also has an Application Behaviour Analysis module. This is a technique employed by CPF to identify malware by analysing their behaviour. This module can identify attempts by certain applications to launch other applications, for example,

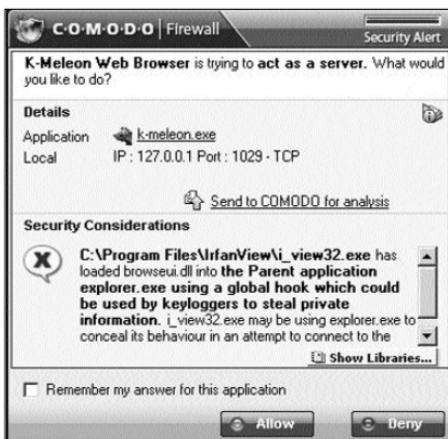
installation routines that launch the browser to take you to a registration page. In such cases, CPF correctly identifies and informs you regarding the Installation routines to access the Net. A similar scenario exists if Trojans launch the e-mail application or the browser to send out data.

CPF can also display the individual data transfers of those applications that have been allowed Net access. This information is useful in arriving at a conclusion about the behaviour of these programs.

The CPF configuration options are grouped under the Security link. Here, under the Application monitor link, all applications and their network access permission can be seen (and modified if necessary). Setting the permission status to "Ask" will make CPF seek confirmation before granting access. You can make the modifications after selecting and clicking the Edit link. The Component monitor lists all the OS components and their permission status. Only you, the user, can change the status.

In case an application needs a specific port to be opened for communication for special reasons, that can be done from the Network Control Rules under the Security link. Here, a new rule can be manually added to allow specific ports to be unblocked by CPF.

Under the Advanced link, the Application Behaviour Analysis parameters can be modified. The Advanced Attack Detention and Prevention module can be used to prevent excessive data from being sent. In an attempt to hog a server's resources, a lot of requests can be sent to it, and when this happens, the server becomes sluggish and can even crash. Such attacks are called Denial of Service (DoS) attacks, since the server can't respond to



CPF, like other firewalls, checks outbound data and decides if it is permissible

other requests in time, leading to a “denial of service” for those users who sent out the other requests. CPF can be configured to limit the rate of requests.

Under the Activity link, the real-time status of applications connected to the Net can be seen, and if necessary, the connection can be closed.

5.5.3 The ZoneAlarm Firewall

The free version of ZoneAlarm includes an e-mail scanner. After installation, a configuration routine will help demarcate the Internet zone from the trusted zone.

ZA will alert the user every time an application tries to access the Net or when a remote site tries to access your PC. Compared to CPF, the user receives fewer alerts. ZA does not seem to detect suspicious application behaviour, as in the case when a browser is launched by some other program.

In the ZA interface, under the Firewall link, the level of monitoring of the scan engine can be set. Here, under the Zones tab the different zones can be identified.



With CPF, you get easy control over which ports are open



An overview of ZoneAlarm



ZoneAlarm pops up alerts—like this one, when a program requests access

The Program Control link can be used to set the behaviour of the firewall when a program tries to access the Net. The free version does not allow the highest setting, which can be used in the paid version to check abuse of trusted applications—as with CPF. Under the Programs tab, trustworthy applications that are to be given network access are listed, along with the scope of permission. In case the permissions given to certain applications need to be modified, you can do it here.

An anti-virus monitoring tab allows ZA to work with an installed AV to alert you if the AV has been compromised. On the reviewer's system that had both AVG and Avast!, ZA failed to detect either. And since the AV is capable of alerting the user to newer updates, the point of this module is lost. It's there probably to advertise ZA's own AV product.

Overall, a download size of 40 MB might make it seem like ZoneAlarm is more feature-rich, but it seems to offer much less than does CPF (whose installer is less than half the size).



ZoneAlarm also pops up alerts like this one (when data tries to get into your PC)

5.6 Anti-spyware

5.6.1 Introduction

Spyware started off as benign programs that reported back to a central server anonymous statistics about the user that helped the server create more effective ads. The same tactic soon began to be used for nefarious purposes. Free screensavers, weather information programs, and such are the usual suspects here. It needs to be noted that most anti-spyware are too focused on protecting every cracker's favourite whipping boy—IE. So, anyone sparingly using IE will not find the same use for these applications. Conversely, anyone wanting to avoid infection from spyware should consider using other browsers, like Opera or Firefox.

5.6.2 Spybot—Search & Destroy

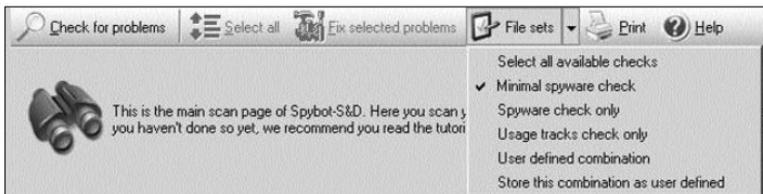
One of the best tools for countering spyware; besides cleaning the system of unwanted applications that have already gained entry, it can also take preventive measures, and includes a on-access scan similar to AV applications.

After installation, Spybot—Search & Destroy (SSD) downloads the signature files that are needed to recognise spyware. Without this, SSD is useless, so the user needs to be online to complete the installation of SSD.

SSD can scan a system for spyware under the Search & Destroy button in the main window. This is a lengthy process. After the scan, objectionable entries are displayed, and you can choose to fix selected entries. The scope of the scan can be modified using the File Sets button. Here, you can set SSD to check for usage tracks, or run a minimal spyware scan—which is considerably faster.



A typical SSD results window



The scan options in Spybot–Search & Destroy

Using the Recovery button, you can undo the Registry changes that SSD performed as part of the cleaning process. This is needed only in case the cleaning renders certain applications non-functional.

Immunisation is a proactive step to ensure that the system is insulated from attacks. SSD can immunise a system by using the controls that already exist—like changing the Hosts file, or setting the browser to prevent certain sites from sending data. The options to do so are presented under the Immunize tab. You can select and apply the suggested immunisation measures.

“Teatimer” is the real-time scan agent in SSD. This monitors the Registry and informs the user about any application’s attempt to modify the contents of the Registry.

In the default operating mode, SSD doesn’t offer any configuration options. This can be changed by going over to the Advanced mode of operation. Under the Advanced mode, the Tools link



The “Recovery” options in SSD

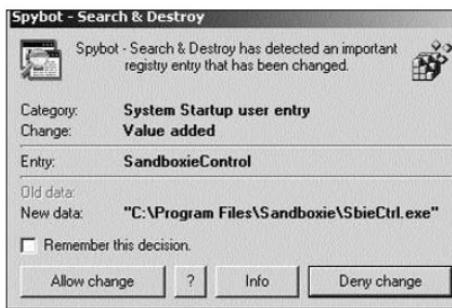


SSD suggests immunisation measures

offers additional tools to combat spyware. The modules under the tools link can be activated or deactivated.

Tools are categorised under different headings like BHOs (Browser Helper Objects), another name for plugins, and can be used to clean unwanted BHOs from IE. System Startup lists those applications that will start along with the system; selected applications can be prevented from loading here. The Process list shows all processes running on the system, and allows the user to terminate selected ones. There are other mod-

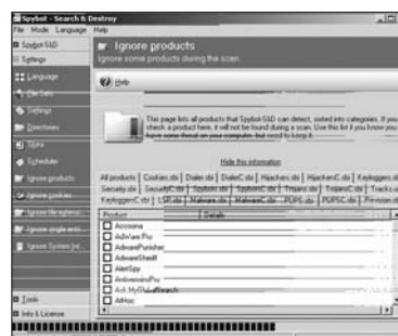
Under the Settings link, various aspects of SSD can be configured. The Filesets link allows the user to select which filters to use for scanning for spyware. The Ignore Products and Ignore Cookies links allow the user to prevent SSD from blacklisting and removing certain entries which may be needed for the proper functioning of installed applications. The Settings link allows the user to change the interface aspects of the program.



SSD can warn you if an application tries to modify the Registry



The Modules under the Tools link



A product checklist for Spybot—Search & Destroy

5.6.3 Hijack This

An advanced anti-spyware tool, you'd only need it if your resident anti-spyware application fails to rectify a "situation." HT includes many tools; it is a Registry checker / editor, it can check the start-up files list, it can delete locked files at boot-time. HT was an independent entity till all the good press made it a target for acquisition by Trend Micro, an anti-virus manufacturer. But it retains its name and functionality.

For best results, to ensure that spyware can be easily deleted, it is recommended to boot the system in safe mode (by pressing [F8] while booting) before running HT.

When HT starts up, you are presented with a window that offers to perform a scan. On running the scan, all Registry entries that are deemed harmful are listed. Selecting an entry and clicking the "Info on the Selected item" reveals additional information about the object and the proposed action HT will take to counter it. You can select and delete entries from the list. HT makes a backup of the deletions, so if needed, any of the deleted Registry keys can be restored. HT can also create a log file of the scan, which can prove useful if you'd like others to suggest which entries to mark for deletion.

The Config button reveals additional tools. "Delete a File on Reboot" can be used to delete a file that is presently

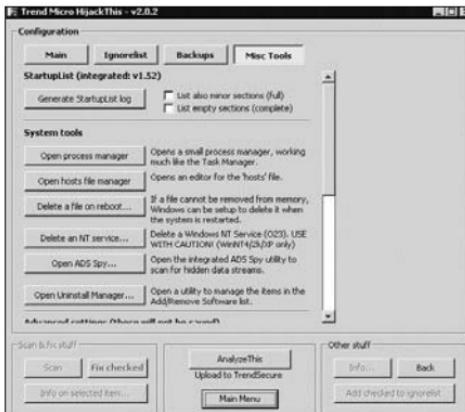
being locked by a process and hence cannot be deleted. Malware applications are usually launched at boot-time and remain in memory; this makes it difficult to delete the file, and stopping the particular process is fruitless since they are restarted almost immediately. This tool will be of use in such cases. The Uninstall Manager can be used to delete entries from the Add / Remove Programs applet. The Main button opens the settings dialog where HT's behaviour can be modified; the start page for IE can also be



A typical scan results window, with info on one entry

set here. The Backups button lists the Registry modifications made, and a selected modification can be undone with a mouse-click.

Post buyout by Trend Micro, the Analyze button made an appearance. This allows Scan results to be uploaded to the Trend Micro site for analysis.



The Miscellaneous Tools section of Hijack This

5.7 Anti-rootkits

5.7.1 Introduction

The hullabaloo over Sony's use of rootkits highlighted the deficiency of most anti-virus applications of the day. The rootkit escaped detection and needed special tools. The danger that Sony's rootkit could be cracked and misused (by malicious people) to insert programs that dodged scans by the latest anti-viruses was quite real, causing Sony to discontinue the practice.

5.7.2 Threatfire

While AV products require an updated pattern file in order to be effective, some applications use behaviour analysis to identify malware. These applications, because of their different hunting skills, remain effective without frequent updates, and can therefore supplement the AV.

Suspicious application behaviour includes autonomously trying to access the Net, autonomously modifying the contents of the Registry, and autonomously creating files. Applications that exhibit such behaviour can be flagged down without the need for a pattern file. You can then explicitly permit or disallow the application.

Threatfire (TF) is a behaviour-based anti-malware application, though it also performs signature based scans. After installation, TF can perform a scan for malware on the system based on the already-included pattern file. The free version can only identify rootkits during such a scan, but TF also has a real-time module that monitors for all types of malware, not just rootkits.

All applications that have been permitted or blocked are displayed in the Threat control centre, where, if needed, a decision can be revoked by the user.

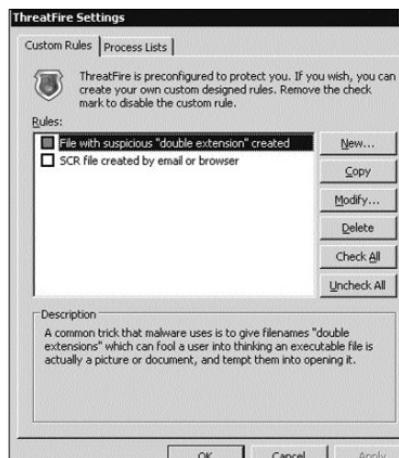
Under the Advanced Rules link, the user can also create custom rules which will be used by TF to identify suspicious applications. These are in addition to the default behaviour rules.

5.7.3 Panda Anti-rootkit

This is one of the many dedicated rootkit hunters that have been released by manufacturers of other security products. Panda also makes anti-virus software. This is a one-trick pony that scans for rootkits and cleans all entities related to it.



The scan types in Threatfire



The results of a Panda Anti-rootkit scan

5.8 History Cleaners

5.8.1 Introduction

Any browsing activity leaves a trail on the system. Temporary file folders, browser history, browser address bar, browser profile assistant, and the Registry are some places where it is possible to find information that can be used to piece together a user's browsing history. History cleaners are a group of applications that can delete the relevant information from these locations and make it difficult to find where the user has been. On a "personal" computer, with all security paraphernalia installed, a user can remain relatively protected. But when the need to access the Net over a public PC arises, the following tools can become imperative.

5.8.2 CCleaner

This is comprehensive history cleaner. It can clean the trails of all three major browsers—Opera, IE, and Firefox. And while it is at it, it can also clean the usage history of commonly-used applications like Office 2003, WinRAR, Adobe Flash Player, and Windows Media Player. It can also erase many other chunks of data that are no longer relevant—like broken shortcut link and Registry entries, thus contributing to a leaner system.

By default, CCleaner looks at all the locations where browsers and the OS leave information. The user can select which areas to ignore, if needed.

Under the Registry and Tools section, there are additional cleaning options available. Under the Options section, the user can manually specify any other folder or file that CCleaner should erase along with the rest.



The results of a CCleaner scan

Locations to ignore can also be specified. Since simple deletion does not physically remove the data from the hard disk (allowing it to be easily recovered), CCleaner can be made to overwrite the deleted data (up to 35 times) for greater security.



The options you get in CCleaner

5.8.3 Sandboxie

Picture this: you come across a program that promises to make your life easier, either by speeding up your system, or automating certain tasks, or increasing your Net speeds—and it costs nothing. You download it and install it without paying any attention to the EULA (End User License Agreement). It then asks you to restart your PC, which you eagerly do. After the restart, you notice that the boot time has increased, opening programs takes eternity, and your CPU is always running at 100 per cent. That's not all; your browser seems to be showing a few additional buttons and controls, the home page has changed, and your System Tray has a few more icons. You decide to uninstall the program, but are asked to connect to the Net to download the uninstaller. Irritation? Anger? *Déjà vu?*

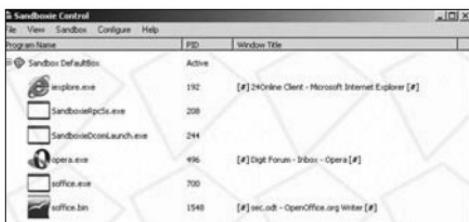
You could restore the system to an earlier date if you'd recently created a Restore Point (or hadn't stopped the Restore service to make room for the latest FPS). But if you were that cautious, you would have read the EULA which, somewhere near the bottom, mentions that the program is ad-supported. You would also have asked around in a few newsgroups / forums for reviews about the program before installing it. Of course, after the incident, you will find many online sources telling you that the program is a front for spyware / malware.

For moments like this, you need a Sandbox. As kids, a sandbox was a great place to do stuff that you wouldn't do over hard

ground out of fear of injury. In software parlance, a sandbox broadly refers to an isolated environment in which a user can perform an action without fear of affecting the rest of the system.

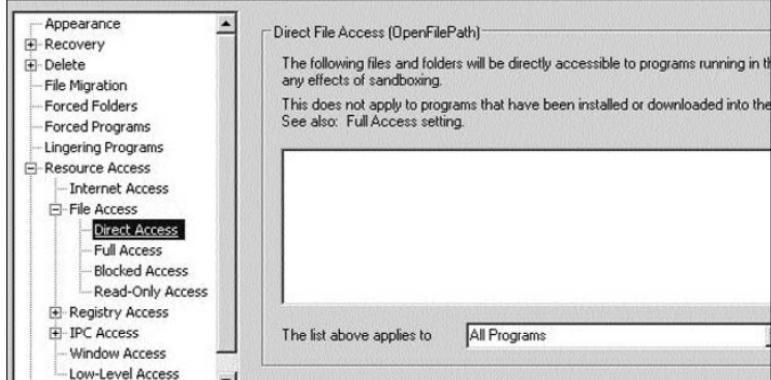
Such an environment is needed when dealing with new programs, since it is possible that the user may not be able to foresee the changes that the program will make to the system. And there is always the probability that the uninstallation routine may not thoroughly remove all remnants of the program from the system.

Sandboxie is a software that creates a sandbox environment on a user's system. It creates a layer between the applications and the user data. After installation, Sandboxie creates a file in its program folder called a Container, which is used to mimic the program paths on the hard disk. Sandboxie allows reads from the hard disk, but doesn't allow writes without express permission from the user. This way, it prevents any direct manipulation of existing data. Deleting the Container removes all remnants of the activities in the sandbox. Sandboxie allows programs of any size to be installed—the free space on the hard disk is the only limitation.



The controls in Sandboxie

Sandbox Settings - DefaultBox



The settings available in Sandboxie

The user can also run, in the sandbox environment, any application already installed on the system. This ensures that any changes during that session are not stored on the system. This is quite useful when browsing the Net on a public PC (Sandboxie needs to be installed first, of course). Any tracks that are usually stored by the browser on the system, which can reveal your browsing habits, are restricted to the sandbox. Deleting the sandbox ensures complete privacy. Any new programs can also be installed in the sandbox environment without causing any changes to the host system.



"Recovery" in Sandboxie means saving data to disk

Sandboxie integrates with the system and is easily available from the context menu. Any application can be run sandboxed in a jiffy. Windows of sandboxed apps have their title boxed between a set of characters, by default two of these: [#]

Sandboxie allows programs running in the sandbox to make changes to the host system, for example to save a file, only with the express consent of the user. Data that the application wants to write to the hard disk remains in the Sandbox container. Sandboxie calls the process of saving this data to the host hard disk "recovery." Sandboxie can be set to inform the user every time there is data to recover, or only after the sandboxed application had been exited. Files that have been created by the application in the sandbox mode can be viewed by changing the viewing mode to Files and Folders. In this view, individual files/folders can be recovered to a user-defined location.

5.8.4 Mojopac

Mojopac is a tool that operates on similar lines to those of Sandboxie. It is intended to be a portable Desktop that can be carried around on a Flash drive. After it has been installed on a Flash (or any external) drive, on activation, it presents a second Desktop environment. Here, applications can be installed with-

out affecting the host system. And once all the applications are installed and configured to match the user's preferences, plugging the external drive into *any* system will present the same, familiar computing environment. From a security point of view, the ability to create a secure environment in Mojopac (with all the security tools installed) and using it irrespective of the infected state of the host system holds a lot of promise.

However, on the two systems that this reviewer tried Mojopac, the host system was affected by the modifications made in the Mojopac environment. For example, installing a browser in Mojopac resulted in the same application being installed on the host system, with just the link showing up in the Mojopac Desktop. Besides that, Mojopac had the irritating habit of preventing any host control from working when it was booting up—which would take a long time if a Flash drive were being used. The user, therefore, cannot switch to the host system ([Alt] + [Tab]), and is only left with the option of thumb-twiddling while Mojopac loads.

A few wrinkles, it is clear, need ironing out. For the adventurous reader, the application is available at mojopac.com. A word of advice: run it sandboxed.



A Mojopac Desktop

5.9 Anti-phishing Tools

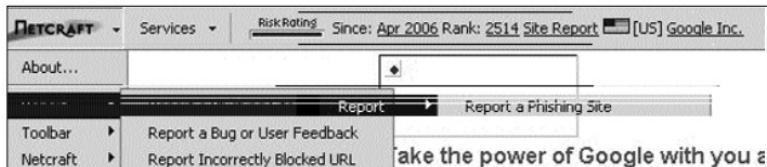
5.9.1 Introduction

Since versions Opera 8, Firefox 2, and IE7, all three browsers have included controls to identify if a certain Web site is “phishy.” All anti-phishing tools work in a similar fashion: any URL the user visits is passed to a central Web site that maintains a list of sites that have been reported to be phishy. If the URL is indeed phishy, the browser indicates this.

To augment the prowess of the browser, additional controls can be installed. These are usually in the form of plugins or add-ons that show up as browser toolbars or buttons.

5.9.2 Netcraft Toolbar

This is available for Firefox and IE. In Firefox, this is installed as a toolbar that displays information about Web sites. By using the Report option under the Netcraft menu, a user can report a suspicious site to Netcraft for further analysis. A wrongly-identified site can also be reported.



The Netcraft toolbar on Firefox

When the user tries to visit a site that has been identified as a phishing site, an alert is displayed. The user still has the option to visit the site. A similar function is available while using the EarthLink toolbar (www.earthlink.net/software/free/toolbar/) and Google toolbar. But both these offer a lot more than what was asked for, and are larger in size. For someone looking for just an anti-phishing tool, the Netcraft toolbar is adequate. Opera users will need the Google toolbar, since the rest do not support Opera.



A sample Netcraft phishing alert

5.10 Miscellaneous Tools

The tools mentioned here add additional layers of security to your system.

5.10.1 Anonymisers

5.10.1.1 Introduction

“Anonymous Browsing” is an oxymoron. Any system on the Net is uniquely identified; this is a prerequisite, since otherwise there cannot be corruption-free data transfer. Without the destination address properly identified, data packets cannot reach the intended system, period.

What a person who is paranoid about privacy can do is ensure that the data path between the source and destination is not direct. A direct path makes the process of identification easy, as easy as a whois lookup (at ip-address.domaintools.com, for instance). By taking a circuitous route, and encrypting the data in transit, the person can achieve a level of privacy that can be invaded only by employing significant resources (it can still be done). It needs to be pointed out that even the “direct” path involves different routers that form the backbone of the Internet. Run a “tracert xyz.abc.com” command at a command prompt to find the number of nodes that route the traffic between you and xyz.abc.com.

5.10.1.2 TOR (The Onion Router)

This application channels data through a circuitous route. The TOR network consists of a network of volunteers who share their systems and bandwidths for channelling data from the source to the destination of any user connected to it. The core TOR application encrypts data at the source, and sends the data packets through a path that passes through many intermediate nodes, with each node reworking the data packet. The number of nodes is usually three. The data packet, as it leaves the source, only has the destination address of the first node in the path. After the packet reaches the first node, its destination address is

modified to reflect that of the next node in the path. This way, each intermediary can only discover (by using the necessary tools) the destination of the packet, not its source. Only at the final node are all the packets unencrypted and put together and sent to the destination.

The TOR application is bundled with a few other tools that make it easier to use it. Besides the core TOR program, the bundle consists of a proxy server (Privoxy), a GUI to control TOR (Vidalia), and a Firefox plugin (Torbutton) which allows Firefox to work with TOR. Running TOR is just the first step: to use it, the browser and every communication tool needs to be configured to send data through it. This can be done by changing the network/connection settings to use a proxy server. This needs to be set to IP 127.0.0.1 and port 8118.

Once this is done, the user should test if the network is doing what it is supposed to do by visiting any site that offers information about the external IP of the user (for example, ip-address.domain-tools.com). The IP reported should be different from the value that is visible by running the “ipconfig” command at a command prompt. In fact, it will be the same as one of the intermediaries in the path. Alternatively, visiting <http://torcheck.xenobite.eu/index.php> will give you specific information about how the TOR installation is performing. The Vidalia control panel can be used to monitor and configure TOR.

The View the Network Button opens a window listing the nodes in TOR in the left frame. The active paths available for use by the user are listed in the middle frame, and clicking on any of these paths will reveal the details of the nodes in them in the right path. The world map will show the location of the nodes. To see the network in action, try accessing any Web site from the browser and switch to the Network map window—and watch the middle frame.

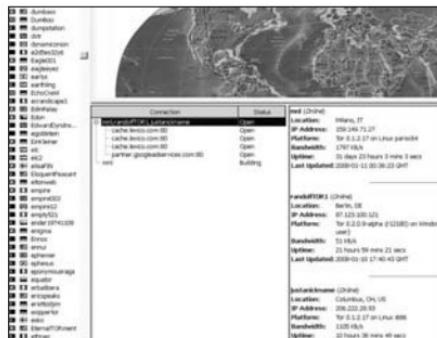
The Bandwidth graph button shows bandwidth statistics in real-time.



Checking if TOR is doing its job



The control panel of Vidalia, the GUI for TOR



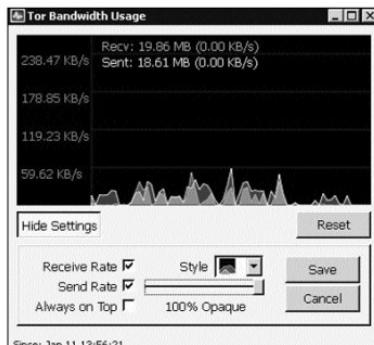
The TOR Network Map view, while accessing dictionary.com

The settings button will open a window to configure TOR. There are very few options that need to be set here.

Those on an unlimited broadband connection can share their bandwidth for the network by clicking on the Setup Relaying button in the control panel. Here, one can set limits to the shared bandwidth, and specify the types of data that are permitted to be channel through. Doing so will cause the username to show up in the left frame of the Network Map window.

Privoxy is a content-filtering proxy server of repute. Since TOR is a SOCKS proxy, the role of Privoxy in the bundle is to act as an intermediate proxy that forwards requests from other protocols, like HTTP, HTTPS, and FTP to TOR.

Note that since the data packets have to make a few more hops to reach the destination, the price of additional privacy is increased network lag. This is not always noticeable, though.



You can see how much bandwidth TOR is using



The main settings window of TOR

5.10.2 Content Filters

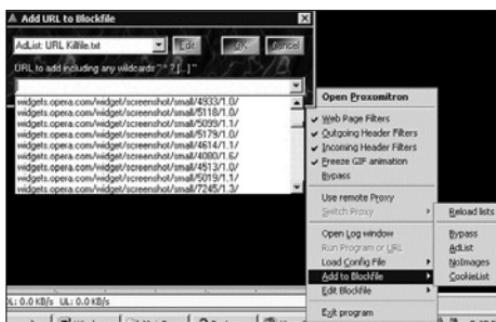
In the normal course of browsing, clicking a link causes a request to be sent to the Web server. The Web server sends the relevant data to the browser. Depending on the page, additional requests will be automatically generated and sent to many other Web sites, like ad servers, for content. Additionally, a browser offers information regarding the user's IP, the Web page previously visited (or referrer), the OS being used, the browser agent, etc. to a Web server, which is (or can be) used to tailor the content sent to the user.

Content Filters act as an intermediary between the browser and the Web server. Depending on the user specification, the application can—in real time—filter the requests sent by the browser to the Web server, thus preventing unwanted Web sites inserting data into what is sent to the user's PC. Content Filters can also prevent the leakage of system information by the browser from the user's system.

5.10.2.1 Proxomitron

This is a content-filtering proxy. Once installed, the browser's connection settings need to be modified to use the proxy. By default, this is IP 127.0.0.1 and port 8080.

Some of Proxomitron's capabilities include blocking of ads; blocking of images; blocking of Web sites; modification or blocking of header information related to the OS, browsing agent, referrer, or cookies; modifying or blocking scripts related to pop-ups, frames, sounds in Web pages, scrolling text, window resizing, preventing right-clicks; and more.



Adding a URL to the Proxomitron blocked list

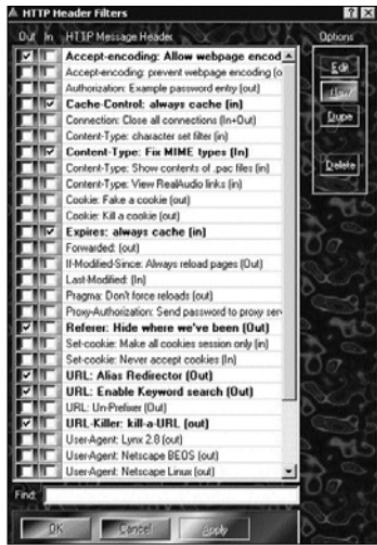
Proxomitron comes pre-loaded with a huge list of sites that have a history of sending spyware, so the user is freed from the considerable labour of adding each one to the browser's blocked

list. Needless to say, besides the privacy aspect, Proxomitron also reduces the amount of data transferred and the page load time, especially in case a page displays many ads from different ad servers.

If the user so desires, additional filters can be easily added by accessing the context menu of the program's System Tray icon. The window will list the URLs recently requested by the browser. A filter can include a URL, or its scope can be increased by using wild characters. For example, the filter blocking "xyz.abc.com/ads" will not block ads from "pqr.abc.com/ads", so the filter can be modified to block "* .abc.com/ads" to block all sub-domains from sending ads. Though the example here refers to ads, the principle for blocking any URL is the same.

Proxomitron's legendary Web page behaviour modifying capabilities can be modified under the Edit filters field in the program window, the colours of which are beyond words.

In conclusion, Proxomitron offers fine-grained control over the kind of sites gaining access into the system through the browser. While individual sites can be blocked in individual browsers, the centralised way of using a proxy server is more effective.



A view of the header filters on Proxomitron



Web page filters in Proxomitron

E-mail Clients



There are two intended audiences for this chapter. If the term “e-mail client” doesn’t ring a bell, you haven’t been online for long; all of this chapter should be useful. On the other hand, if you know that “e-mail client” means Outlook or OE, then read on for an introduction to a few alternatives. (And if you used pine in 1996, don’t send us hate-mail for having included this chapter... there *are* newbies around, you know!)

6.1 The Various Ways Of Using E-mail

Barring a few exceptions, everyone gets their mail in one of three ways: using Webmail, using IMAP on an e-mail client, or using POP3 on an e-mail client. (Let's just call an e-mail client a "mail program"; it sounds less officious).

"Webmail" refers to these: Gmail, Hotmail, Yahoo! Mail, Rediffmail, and more. In this method, you log on to a Web site with your username and password, and you're taken to a page where you can click on links to read your mail. You can click on a "Compose" or "Write" button to write a mail, then on a "Send" button to send it. All this is happening on the Web. You're using your *browser*, the same program you use to visit any other Web site; your browser could be Internet Explorer, Firefox, Opera, or something else.

A mail program refers to these: Outlook Express (it's installed on all non-Vista Windows computers), Outlook (which comes with Microsoft Office), Thunderbird (which was developed by the same group that made Firefox), and others. These are programs dedicated to the sending and receiving of mail; they are not browsers (though some do have browsing capabilities, but let that not confuse you). When you send and receive mail using these, you are not on the Web, meaning you aren't visiting a Web site. And now, there are two main "ways" you can use a mail program—IMAP and POP3.

Two sentences of explanation are needed here: when someone sends you an e-mail, it gets stored on a *mail server*. What you're doing when you use a mail program is, you're retrieving your mail from the mail server onto your personal computer, either permanently or temporarily. IMAP and POP3 happen to be the two ways (mainly) of doing that.

6.2 IMAP Vs. POP3

Now, let's look at IMAP vs. POP3 (forget about the expansions; no-one will ask you). In simple terms, in IMAP, your mail program manipulates the mails on the mail server. Most actions happen on the server; it's like you issuing commands for things to happen remotely (on the server). In contrast, in POP3, you *download* your mails once and for all. When you click a button to receive your mails, they come to your computer, and are gone from the server. Your mails are now like pieces of data on your computer, which you can manipulate (on your computer) using your mail program.

That's the technical difference; what does it mean for you? Well, take two cases: person J, a frequent traveller, who uses IMAP, and person Q, a student who uses only one computer (at home), and uses POP3.

J is using his laptop on a journey. He connects to the Net using a wireless connection, and is billed according to data transfer. He fires up his mail program and tells it to fetch his mail folder list (which is possible only because he's using IMAP). He looks through the list and sees there are fifteen new messages in the Inbox. He then gets the message *headers* (information about the message) only, not the messages in their entirety. This, too, is possible only with IMAP. He sees that someone has sent him two 12 MB mails; they are the same mail, and the sender sent it twice by mistake. He deletes one of them without even opening it (IMAP allows this). He downloads the other.

Then, J notices that ten of the fifteen mails are all from a particular client, and he doesn't need to read them now. He creates a folder on the server called "Some Client", and moves those ten messages to that folder (you need IMAP to do this).

Two days later, J logs on to the Net from a cyber-café. He double-clicks the mail program icon, goes to fetch his mail, goes directly to the "Some Client" folder, and downloads them all.

See what's going on? Everything is happening on the server. The mail program is being used to fire commands at the server, as it were. Nothing is deleted from the server unless (a) J actually issues a command to delete them, or (b) if J marks a message for

deletion, and the server is configured such that it deletes such messages at regular intervals.

If J regularly received hundreds of mails a day, he could have issued a command to view messages that met certain criteria (such as “smaller than 10 KB). Messages on the server can be marked with status flags such as “answered”.

In contrast, Q returns home from a busy day and turns on his PC. He clicks on “Receive Mail” in his mail program. All his mails come to his hard disk, and are gone from the Internet. Now, if Q goes to a cyber-café and opens up a similar mail program to fetch his mail, nothing will show up.

That’s the typical scenario, but actually, POP3 does allow you to keep messages on the server—it’s just a simple setting in the mail program. So, for example, Q might open the mail program at a cyber-café and set it to leave messages on the server. He’d then download his messages, read them, then delete them (so that the next user at his seat won’t get to read them). Then, he goes home, where his mail program is *not* set to leave messages on the server; when he downloads the mail there, it’ll be saved to his hard disk and erased from the server.

So, then, which do you like: IMAP or POP3 or Webmail?

6.3 The Choices

The answer to that question depends on several factors. Ask yourself the following, not in order of importance:

- How fast is my Internet connection?
- On what basis do I pay for my Net connection?
- Do I want to read and write mails when not connected?
- On how many computers do I check my mail?
- Is there a limit on the storage at the mail server?
- Do I have the habit of making backups?
- Is it possible for others to get at my mail?
- Does it feel better reading mail without external content garlanding it?
- Do I have a choice between POP3 and IMAP with my ISP?
- Do I have a choice between POP3 and Webmail at work?

OK, so that's 10 questions already—and a walkthrough of those ten should clear things up a bit.

6.3.1 Connection Speed

Remember that Webmail (such as Gmail and Yahoo! Mail) has several features that take time to download, ads that take up bandwidth, and more. So if you're on broadband, you're fine, but with dialup, the Webmail experience can sometimes be frustrating. Also remember that POP3 can often be a bad thing if you have a slow connection; if someone sends you a mail with a 7 MB attachment, you'll have no way of knowing—you might sit a long while downloading the mail only to discover it's a silly forward. IMAP could be the best answer for those on dialup—you just download the header information, which happens in a jiffy, and download only the mails you want to.

6.3.2 The Money

If you pay for data transferred, POP3 (and, to a lesser extent, Webmail) could be bad news; IMAP should be preferred for the reasons mentioned in §8.3.1. If you have a broadband connection and pay by the hour, even a 7 MB mail won't take long to download—so POP3 is not so much of an offender. Webmail could be a bad choice if you pay by the hour; each transaction (delete mail, read mail, move mail) takes time.

6.3.3 Offline Activity

It should be obvious: if you're using Webmail, you need to be connected to read your mail. On the other hand, if you're using a mail program, you can read yesterday's mail anytime today (using POP3, and with IMAP for the messages you downloaded).

6.3.4 Mobility

Again, it should also be clear that IMAP is to be preferred if you connect from several locations. You treat your mailbox not as a local repository but as a remote one, like we said. Also, Webmail is good if you access your mail from several locations, but you need to remember that load times are slower, so it also depends on the connection speeds at the locations you connect from.

6.3.5 Storage Limits

This is fast becoming a non-issue, but some providers do have tight limits (as low as 50 MB in some cases) on how much mail you can keep stored. In such a case, for example, using IMAP to keep your mail organised on the server would be futile; you'd need to keep deleting mail (or downloading them, as in POP3).

6.3.6 Backups

With a mail program using POP3, your mail is on your hard disk, nowhere else. So you should be in the habit of making backups at least once in a while. With Webmail, this isn't an issue—almost never have people lost mail online. With IMAP, it would depend on the storage limit—if it's large enough, it's like the situation with Webmail.

6.3.7 Privacy

Just remember that with Webmail, others could get at your mail if you don't sign out properly. Remember to uncheck the box that says "Remember me on this computer" (or something to that effect) if you're using a public computer.

6.3.8 The Cosmetics

People used to mail programs feel uncomfortable using Webmail; there's your mail, in the middle of the screen, and to the top and the bottom are ads, and to the left and right are other Web site elements. With a mail program (whether you're using IMAP or POP3), you have a neat window with just the text of the mail. It can feel more personal.

6.3.9 Do I Have A Choice?

At work, your mail program has probably been set up for you; it's unlikely that you'll have the choice of POP3 or IMAP. On the other hand, many ISPs offer you both. For example, Sify does. To find out if your ISP does, just Google "ISP_Name imap" and "ISP_Name pop3" and you'll have your answer! The settings for both (POP3 and IMAP) that you'll need to enter into your mail program for it to work will be on pages at your ISP's site; follow the instructions carefully.

6.3.10 I Love Gmail; Can I Use It At Work?

The idea is, some offices and colleges block Webmail sites. You'll find out soon enough. Even otherwise, it doesn't "look good" if you're often seen with a Gmail window open (and since Gmail integrates a chat facility, you might be tempted to visit www.gmail.com quite often). But remember that Gmail (for one) offers you the facility of downloading your mail via POP3 and/or IMAP to your mail program: in Gmail, just click on Settings, then go to the "Forwarding and POP3/IMAP" tab.

The choice is getting a bit hazy. For one, connection speeds are improving, so much of the virtue of IMAP is getting eroded. Second, many Webmail sites (notably the new Beta of Yahoo! Mail) are giving you the look and feel of a mail program. Third, you're more likely than ever to see Webmail interfaces for, say, your office mail, or your domain mail. And then, you have the fact that, for example, Gmail as mentioned above, gives you free POP3 and IMAP mail (God alone knows why they do it; you don't see any Google ads in your mail program when you use POP3 or IMAP for your Gmail).

In any case, that's 10 questions answered; you decide... and if you decide that you need a mail program (or an e-mail client, if you will), here are a few that deserve mention.

6.4 Outlook Express (OE)

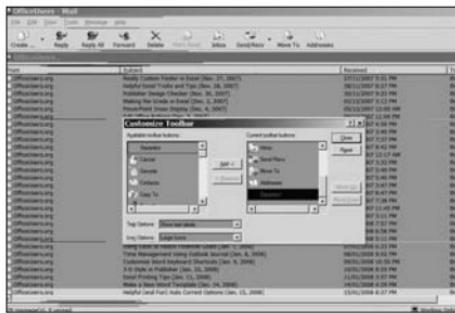
OE must be mentioned because it's there on all Windows computers up until XP. Several people advocate dumping OE, primarily because of security concerns; some have reported data loss when OE "compacts" messages automatically. Then, there are some weird bugs that Microsoft refuses to iron out: for example, when a message is open in its own window, you can delete it from the main OE window... and then when you click Delete on the *message* window, the zombie of a thing just says "The message could not be deleted." Other irritants include, but are not limited to, composition of messages in HTML not behaving as it should, and weird behaviour upon pasting of text into mail composition windows.

Still, it's there...

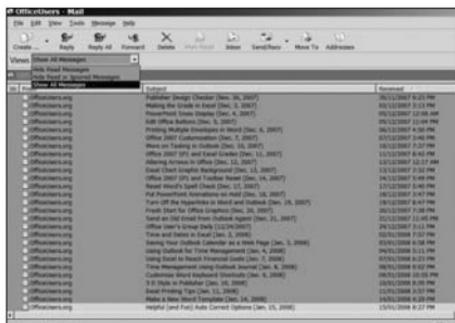
OE is quite simple, and most people could just go ahead and use it. Still, here are a few usage tips.

At the top, but below the line of main menus (such as File and View), you'll see the toolbar buttons. Right-click anywhere to the right of the last button, and you'll get a menu box as in the following screenshot. As you can see, you can customise what buttons should be there on the toolbar, and doing that is as simple as it gets.

In the same area, you can right-click and check "Views Bar". This will bring up a bar that allows you to choose between three views: Hide Read Messages, Hide Read or Ignored Messages, and Show All Messages. The first and third ones are obvious; the second also hides "ignored" messages. Well, it's like this: under the Message menu, you can choose to "Watch" a conversation or to "Ignore" a conversation (or neither, of course). The effect of the latter is that in the second Views option as above, the message won't show up; the effect of the former is that a "Watched" message, and all replies to it, get marked red by default. (For example, you get a mail from Q with the subject "Roses are red"; you reply to it; Q then replies



Customise the toolbar so you only see the buttons you need



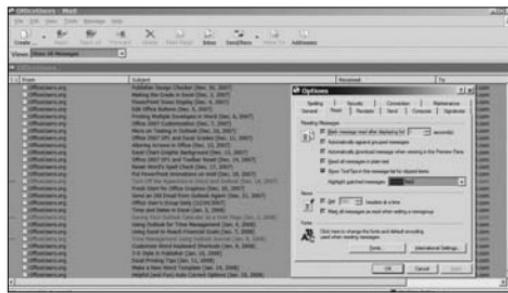
With the (minimalistic) Views bar on, you can choose to hide read and ignored messages

with the subject “Re. re. Roses are red”; both these mails from Q will be marked red.) You can change this colour by going to Tools > Options > Read and looking near where it’s really red.

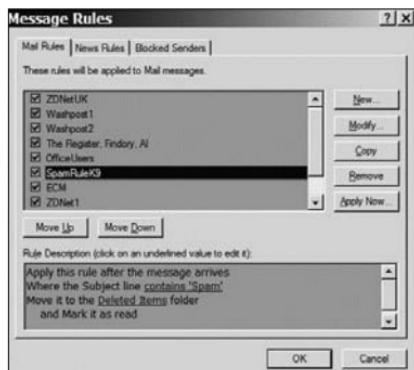
Here’s an important aspect of OE: you can add “Identities”, so different people using the same computer can use OE, each with a different password. The sad news is that the data isn’t encrypted, so if you go to where the messages are stored, you can view all the messages in a text editor—somewhat messed up, but intelligible nevertheless.

Anyway, you can go to File > Identities > Add New Identity and follow the dialogs. After you do this, each time you start OE, you’ll need to select an identity and supply a password. Probably the best use for this (since the messages aren’t encrypted) is to better organise your mail—one identity for work-related mail and one for personal mail, for example.

Then, there are message rules. Go to Tools > Message Rules > Mail. You’ll get a dialog as below, only it won’t be populated. Click New; you’ll get a series of conditions in the top half of the window—select from one of them, then choose, from the bottom half of the window, what to do when that condition is satisfied. As you keep doing this, the rule gets built up. Space does not permit a full description of the types



Here, some messages are being Watched, so they’re in a different colour—which can be changed as described

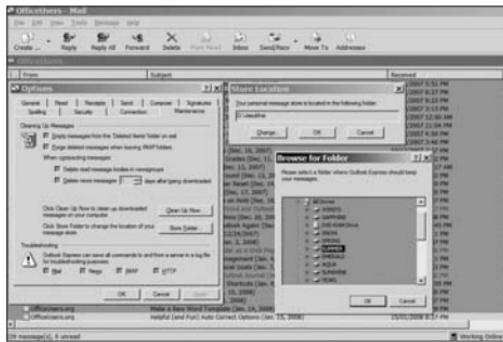


The first dialog in the process of creating a new message rule

of rules that can be applied, but suffice it to say that as you'll discover, rules can be used to auto-move messages when they arrive, auto-delete messages, check for messages with a certain word in the subject line and auto-forward it, etc.

Tools > Accounts > Mail is where you'll be going to add a new mail account; the Wizard guides you through the steps, but you need to know all the details, of course.

And finally, for the all-important Store Folder: OE stores your mails in an obscure location. Since you should be taking regular backups, it would help if you changed it to someplace you know.



Change the Store Folder in OE and keep backing that folder up

Go to Tools > Options > Maintenance > Store Folder, and the location will be displayed (as in the middle window below). Click Change, and you'll be asked to browse to a nice folder where your messages will be stored, as in the right-most window below.

That's all, folks! There really isn't too much more to OE—all we can say is, look in Tools > Options, and you'll find several tabs, all of which contain rather self-explanatory items. Set them according to what you want, for example, what your HTML mails should look like.

6.5 Windows Live Mail (XP and Vista)

It's pretty, very pretty. It's like Outlook Express, but it's pretty. It can search the full text of e-mails for a keyword you specify in a convenient box on the main window, but for that, you need to download Windows Desktop Search (by clicking on the link right there). Windows Live Mail is, essentially, a good-looking replacement for OE on XP and for Windows Mail on Vista.

First off, you need to have a genuine copy of Windows to be able to install Windows Live Mail; when you go to the download link, <http://get.live.com/wlmail/overview>, and click on the “Get it free” link, you’ll be downloading a small file that will probably ask you to get your computer up to date at Microsoft Update—and the first step to that is, of course, the Windows Genuine Advantage Validation. After that’s done, the downloader gets to work and installs the software which, right on first run without any prompting, brings up the following:



Live Mail will import your OE mails when you install it—without prompting you

So there you have it—all your OE mails imported. The first thing you’ll notice is the way your mails are categorised: they’re under your individual accounts (POP3 or IMAP), each

of which has the “default” set of folders such as Deleted Items, Inbox, etc. This is followed by an automatically-inserted Hotmail, which in turn is followed by any other folders you might have created (such as “Work” or “Personal”.) You’ll get it when you take a look.

Then, you can right-click on a message and select “Junk E-mail”; then, you can perform one of several actions, including marking the mail as junk, or adding the sender’s domain to the Blocked list. Yes, Windows Live Mail has a junk mail filter much like that in Outlook (MS Office)—you can choose from two levels of strictness, or choose to go by lists of blocked and allowed senders. Early reports indicated that the junk mail filter didn’t work well at all, and though we didn’t test the filter, there are so many bad reviews out there that we’re pretty sure it’s bad.

Then, like we said, for some strange reason, while OE would search your mails for a



Note the toolbar at the top, and the different way of categorisation

phrase or word, Windows Live Mail requires you to install Windows Desktop Search (WDS) for the purpose. Not that WDS is a bad thing, but *requiring* you to install it could be.

Do take a look under Tools > Safety Options. There are a lot of options here, including the phishing filter, how suspected junk mail should be treated (go ahead, give it a try!), a list of blocked mail senders, and more.

Most importantly for those who like things like iGoogle (everything in a single place), after you log in to Live from within the client, you can send an IM as a reply to a contact when you receive a mail from that contact. You can call (Live has a VoIP app just like in Yahoo!

Messenger, for example). And you can post the message currently being displayed as a blog post on Windows Live Spaces. Not revolutionary, but many will find it useful. However, this also means that the program is sort of bloated (which is not unjustified, considering all those functions)—so if all you want to do is send and receive mail from a single account, you'll find it more sluggish than necessary.

Finally, to reinforce the “pretty” idea, you have different views. Just right-click somewhere near the top of the window and you’ll get the option to change views; a click here and a click there, and it all looks different. Go ahead—have some fun!



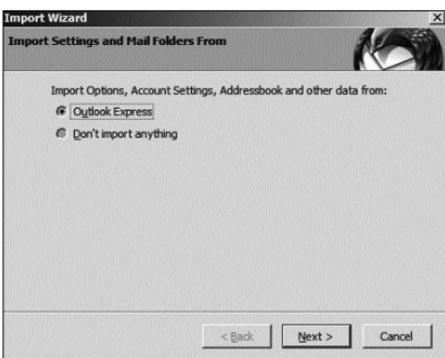
Right-click on a message and give it a bad name



Here's what Microsoft cryptically calls the "Reading Pane" (why, anywhere there's text, you read!)

6.6 Mozilla Thunderbird Version 2

Thunderbird is to e-mail what Firefox is to browsers, but it really isn't—in the sense that it didn't create quite the stir that Firefox did. Still, it's a good e-mail client, period. Now since this is a third-party mail program, you might be a bit worried about some mails remaining in OE and new mails coming into Thunderbird—but upon installation, the first thing Thunderbird does is to offer to import your mails, as shown.



A comforting dialog, this

Now Thunderbird (as opposed to the plain OE or its mostly-cosmetic replacement, Windows Live Mail) is all about features. There are quite a few worth mentioning, and we'll talk about the more important ones.

First, take a look at the compose window. It's feature-rich. Here are all the things you can do with simple clicks on buttons right on that window:

- Bring up your contact list and select one (or more) from there.
- Save the mail as a file.
- If a message in the main window has been highlighted, you can press "Quote" to quote the message—instead of having the mail program set to quote or not quote by default.
- Apply paragraph styles, too, not just fonts.
- Increase and decrease text size easily, bullet and number, insert pictures, and insert an emoticon from a goodish collection!

Little details can sometimes make the difference between an OK and a great product. In Thunderbird, you can colour-code messages as "Important", "For later", etc.—nice! The menu is right on the toolbar. Then, as below, under Tools > Options, you actually have a separate tab for Attachments: what the program should do with them. You can have it automatically save attachments to a particular folder when you say "Save". In addition—and here's the

cool part—Thunderbird can actually automatically save or open attachments of certain types; you can specify what it should do.

Spam-busting is probably what Thunderbird is best-known for. You have a button on the main toolbar that allows you to mark a message as spam; and the spam filter learns, meaning that as time passes and you mark more and more messages as spam, it gets better and better. We haven't tested it, but reports say the spam filter works very well. Phishing protection has been implemented, too—so you get a warning right in your mail program that a message might be a phishing mail.

Like in Windows Live Mail, there's a search toolbar at the top for when you're looking for a particular message, but going beyond, you can save a search as a folder—and running the search again is as simple as clicking on that folder.

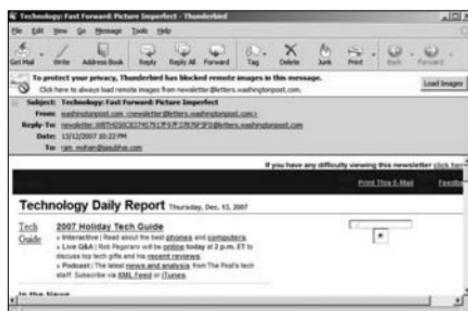
There are also the Back and Forward buttons on the toolbar, which are supposed to make it very convenient to go back and forth between the messages you read—much like in a browser. As in, when you press Back, you go to the message you read just before the current one. However, it's been reported that this feature may not always work correctly. When we tried it out, the buttons just remained greyed out.

Now, as you might know, it's not always safe to download images in e-mails, especially if you don't know the sender. By default, Thunderbird blocks images from loading, and the message window appears like the one alongside.

Now the same thing happens in Outlook Express, and you can click a button to load the images. However, in Thunderbird, you can



Tools > Options is usually boring—not so here



Images are blocked by default in Thunderbird

also, with one click, opt to always download images from that sender.

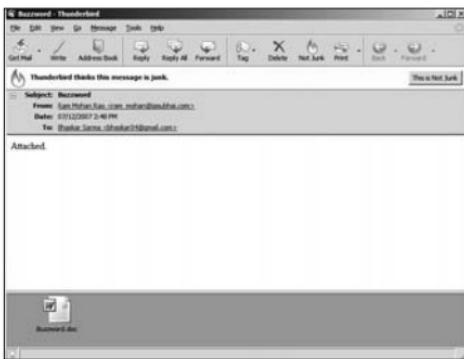
Finally, remember that “training a junk mail filter” isn’t difficult in the least. When Thunderbird is learning, and it thinks a good message is junk, it says it in plain language right in the message:

“Thunderbird thinks this message is junk.” You can then choose to mark it as junk or tell the program it’s not. Also note in the screenshot below how attachments appear as pretty icons beneath the text of the message.

And then, of course, there are the themes and plugins (which Mozilla calls Add-ons). Just visit <https://addons.mozilla.org/en-US/thunderbird/browse/type:1>. Extensions can make the base software much more functional, and it’s also nice to have your program customised your way—and that could be, for some, the biggest reason to switch to Thunderbird. Here’s a quick sampler of the best extensions available:

- Enable bi-directional access to Google Calendar
- Signatures can be a pain in the behind when you decide you want to change it; there’s an extension to quickly switch the signature on/off or choose a new one from a predefined set
- “Did your friend send you a cool link? Don’t want to open up your web browser to view the link? You don’t have to! Just view it in Thunderbird! ThunderBrowse it baby!”
- Extract all attachments from selected messages and then delete, detach, or mark-read

There are about 360 available. Of course, many are crappy, but you’re sure to find, say, 20 that you like... and that’s a lot!



Messages can't get more user-friendly!

6.7 The Bat!

This one isn't freeware, but too many people swear by it—that's why we're mentioning it here. It costs \$35 (Rs 1,400), so let's take a look at whether it's worth that much. The download location is <http://www.ritlabs.com/en/products/thebat/download.php> (The Bat! v3.99.29 Home Edition).

This is one program with a lot of features, and for a full explanation, we'd recommend you visit <http://www.ritlabs.com/en/products/thebat/features.php>. In what follows, we'll only describe the outstanding features.

So well, when you install The Bat!, you get this rather unconventional, professional-looking message:

The message (note that this dialog appears before the program has even started up) says that regular backups are important. It offers to, basically, package



The Bat! means business!

everything in the mail program into one single file, on a regular basis. You know where the file is, so you could copy it to another hard disk or to optical disc when you feel like. If anything goes wrong—you need to reinstall Windows, or your hard disk crashes, or whatever—all you need is that one file to be back in business.

As you'd expect, you can import your mail from OE or from one of several other programs. You'll notice, when you right-click a message, that there is a rich variety of options: deal with the attachments, Park or Un-park it (explained later), assign it a Colour Group (explained later), and more.

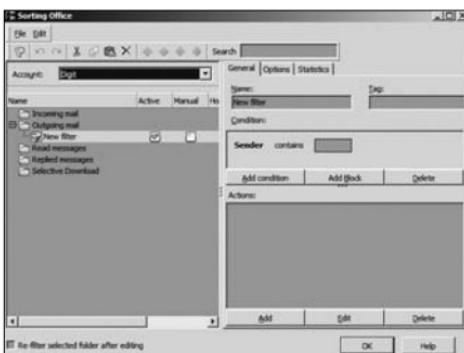
As an example of how comprehensive in terms of mail management features the program is, the filters aren't built up condition-by-condition as in Outlook or OE. It's different, and it's likely

you'll find it more functional. Under Accounts > Sorting Office/Filters, you'll be able to configure the filters for various categories of mail.

Now there are many features unique to The Bat!, so here are some of them, not in order of importance:

1. The Mail Dispatcher

A very useful feature you'll notice is The Bat!'s Mail Dispatcher, using which you can manage mail on the server, even with POP3 accounts. This, of course, is very useful when you receive messages with large attachments and you're unable to download them because you're on dial-up. To "dispatch" your mail on the server, you go to Account > Dispatch mail on server, and choose "All messages" (assuming you've set



In The Bat!, you have a complete filter control centre



Right-clicking on a message gives you a lot of options

messages to be deleted from the server after they're downloaded). The Bat! connects to your mail server, gets the headers, and presents a list of messages with their current flags:

- Read (Mark message as read without downloading)
- Receive (Retrieve message)
- Delete (Delete message without downloading)
- Open (Retrieve message and open it immediately)

You can flag the messages as you'd like; Select Message > Execute and specify what you want.

2. Message Splitting

The Bat! allows you to automatically split large messages into smaller bits. This is very useful to circumvent size limits on messages; however, the recipient needs The Bat! (or an RFC1521, Content-Type: Message/Partial compatible software, as in the help file) to be able to put the message together.

3. Message Encryption

The Bat! helps protect your privacy by employing OpenPGP (built in) to electronically sign or encrypt your e-mail. Refer to the Help file for instructions on how to install PGP support.

The Bat! also offers two S/MIME (a standard for signing and encrypting mail) implementations: its own internal implementation, and Microsoft CryptoAPI.

4. Virtual Folders

Virtual folders are a productivity measure, useful if projects (and more than one e-mail account) are involved. They contain only messages meeting certain criteria, and you define these in the Virtual Folder dialog. Virtual Folders can be within accounts, or they can be common folders outside any account. They only hold references to the original message.

You could use virtual folders to keep track of unread messages in different folders for specific projects, or to keep track of flagged messages in different folders, for example.

5. Colour Groups

The program gives you the capability to define Colour Groups for your messages. They provide a quick way to differentiate between messages, and therefore to categorise and sort them. With a Colour Group, you can define font style, colour, etc. for a message in a Group.

6. Exporting the Address Book...

The Bat! can export your address book into the following formats:

- LDIF
- vCard (electronic business card)

- CSV (comma separated plain text file)
 - Tab Separated File (plain text file)
 - Address List (plain text file)
 - INI file format (plain text file)

7. Message “Parking”

This feature prevents accidental deletion of messages. When a message is “parked,” you won’t be able to move it to another folder, and it will not be automatically deleted or purged—you’ll have to make a deliberate decision to delete it.

And finally, to give you an idea (yet again!) of how serious The Bat! is about its job, take a look at the Customise menu (which you access by right-clicking just below the window bar). You can customise the Menu & Toolbars, Shortcuts, and Options...every little detail.

About the price: you do get a lot, but \$35 is too much, we'd say. The privacy and security features are unbeatable, but how many of us send out extremely confidential stuff via e-mail? Still, if you've got the money to spend...

6.8 IncrediMail Xe

Here's a screenshot to get you into the spirit of things. This is what you'll see when IncrediMail Xe starts off for the first time:

Then, when you, out of curiosity, click on the option that says to turn junk-mail protection On, you get this:

It's about \$25 (Rs 1,000) for a year, so, well, you know you don't need to click on "Click here".



Yes, it does look like it's for kids...

In any case, there are the “Click here” kind of offers at many places, so be warned.

Regarding skins... we don't know many there are, because we didn't click there; still, here's what the skin chooser looks like (it'll probably be populated when you do click there).

A feature that many will find useful is Babylon for Incredimail (Babylon is a pretty good source of cross-lingual information; it includes a thesaurus and dictionary, information about keywords from various sources, and, most importantly, translation). You can select a word and

[**Ctrl**] + right-click, and some Babylon entries for the item will show up. Here's what shows up when we did that after having selected the word “preferences”, with information and translation in German having been selected:

All the rest is pretty of Incredimail is pretty much OE-like. We'd recommend it if:

- You're a kid
- You're a kid at heart (or just like fancy graphics)
- The inbuilt dictionary and translation features are appealing
- You have the money to upgrade by buying new features as we've mentioned (or go the whole hog and just buy Incredimail Premium)

In any case, you don't lose much trying it out... what we suspect, though, is that the aggressive demands for money will initially put you off.



We want your money! Buy our spam filter!



We want your money again! Buy our skins!

Newgroups and RSS



With the ability to share things in real-time, it's no wonder that the Web has revolutionised the way the world shares news and information. RSS readers are perhaps the most popular way of syndicating news that you like and presenting it all to you in a unified interface. Newsgroups on the other hand bring like-minded people together, and allow them to share news of interest. This chapter will focus on both.

7.1 Newsgroups

7.1.1 Introduction

Newgroups (NG), Usenet and BBS are one of the oldest services available on the Internet. They are similar to the forums that are widespread in the Web at present, in that these are places where people come to discuss, mostly through text messages.

The core of Usenet is a network of newsgroup servers. A newsgroup server operates on the NNTP (Network News Transfer Protocol), similar to a Web server running on HTTP. Each server hosts various newsgroups, and the content in all the newsgroup servers is synchronised on a regular basis. A user starts off a discussion by sending a message to the newsgroup address on any NNTP server. Other users connected to the same NNTP address, and tracking the newsgroup can read that message. After the NNTP server syncs with other NNTP servers the message becomes visible to all NNTP users (this process is sometimes also called message propagation). Similarly, replies to the message are also sent by users to the NNTP server they are connected to, which then syncs the data with the rest of the servers. Usually Internet Service Providers host NNTP servers for their subscribers, though such a service doesn't seem to be offered by any of those based in India.

In the original scheme of things there were eight top level categories of newsgroups, namely soc, comp, sci, rec, talk, humanities, news and misc. Each of the main categories caters to specialised topics: "soc" for discussion related to society, politics etc; "comp" for discussions on computers and related technology; "sci" for scientific discussions; "talk" for discussion on general topics; "news" to discuss latest events; "misc" for un categorised topics; and "rec" to discuss hobbies or other recreational activities. The "alt" category was added later on, and with fewer restrictions on creating new NGs in this category, it now has the most NGs. Starting a new NG in the "alt" category is more easier than starting one under any of the rest. (But it is still more difficult than starting your own Web forum). Since there is no central authority ensuring a strict classification for NGs, it is common to find NGs under top level category names

besides the main 9, like “hr” specifically for discussions in Hungarian, etc.

There are subgroups under each of these main categories, and recursively, more subgroups under those. A newsgroup can be uniquely identified by its address which consists of the names of all the groups it is under—for example comp.lang.c++ is a newsgroup discussing the C++ programming language which was classified under the languages newsgroup for discussing languages, under the comp groups for discussion on IT topics. It would be hard to make even a rough estimate on the actual number of newsgroups, but it is fairly correct to state that one exists to cater to nearly every human interest.

When compared with Web forums, newsgroups are comparatively “serious” in nature. Unlike forums with their reps and smiles and emoticons and avatars, newsgroups are predominantly all text. But, there are NGs that contain pictures and other content. In fact, alt.binaries is an NG that is specially used to post large files like movies etc, almost just like P2P networks. It is probably an indication of the age of the users, but the discussions in newsgroups tend to be serious, and the language formal and mature. Unlike forums where there will be a censor or admin to monitor the content, anything goes in most newsgroups. People are free to post any topic in a newsgroup, but if they expect any meaningful discussion it is best to practice self-censorship. Off topic posts usually get a gentle reminder to repost in the appropriate newsgroup or are simply just ignored, just like all the spam that also floods these servers.

Most NNTP servers charge for granting access to users. Some require you to register. But there are others that are free. Since binaries take up a lot of space and bandwidth, it is almost impossible to come across free NNTP servers offering this category of NG. It is important to realise that the content on most NNTP servers are regularly purged. So the chances of finding older topics are quite slim, but this varies from server to server with some having a shorter retention time than others.

To access NGs a NNTP client or newsreader is needed. NGs can be also be accessed from online sites like groups.google.com,

which regularly updates NG data. Google has its own version of newsgroups called Google Groups, which should not be confused with the original NNTP newsgroups. And given Google's never-ending and ever increasing hunger to store and process data, one can be fairly certain that the posts will be there for a long time to come, even after the original NNTP servers have purged them.

Given the similarity between e-mail and newsgroup clients, most e-mail clients can also work as NG readers (and vice versa!), and include the necessary options to configure an NG account. Though, they will lack some of the features present in applications positioned primarily as NG readers.

Configuring an NG client is similar to configuring an e-mail client, but much simpler. One only needs a newsgroup IP to begin with. If the NG only allows registered users, a username and password will also be needed. A list of free NNTP servers can be obtained from online sites like http://www.dmoz.org/Computers/Usenet/Public_News_Servers/

7.1.2 Terminology

Cross Posting: This refers to the practice of posting the same message to different newsgroups.

Binaries: These refer to the non text part of a message—usually pictures, video files, audio files- that are attached with the message. alt.binaries is an NG that is dedicated to posting messages with binary content.

Retention Time: NNTP servers do not store topics indefinitely. The duration for which they retain the topics is the retention time. This varies from server to server.

Multi-part Message: This applies to messages with large attachments. Since servers have limits on the file attachment sizes, users split the attachment over more than one message. These messages constitute a multi-part message. Newsreaders that support multi-part messages are capable of identifying the different parts of a message and joining the attachment to recreate the complete file.

Message Encoding: This is a issue that is applicable more to binary messages. Since NNTP servers were designed to handle text data, to be able to pass binary data through them the message has

to be encoded. Different encoding schemes exist, like yEnc, MIME and UUE. A user intending to download binaries should choose a client capable of handling files encoded in either of these schemes. yEnc is the most frequently used scheme in alt.binaries.

Header and Message Body: A header is that part that carries meta-information—like NG address, poster address, subject, encoding scheme etc. The message body carries the actual message or attachment. Most clients can download just the headers of messages allowing you to judge whether to download the message body as well. This saves a lot of time and bandwidth.

7.1.3 Newsreaders

A few Newsreaders are discussed next.

7.1.3.1 Xana News

To use it, under the Accounts menu use the New Account link. Here you can enter details like screen name and the NNTP server address. After the necessary details are entered, the server shows up on the left frame. The next step is to get the list of newsgroups available on the server. This can be done by clicking on the “Show Newsgroup List” after selecting the server. Depending on the number of NGs in the server, this could take a long time, but this needs to be done only once (or rarely). The list of all the NGs on the server is displayed.

To quickly find an NG on a certain topic the Filter box can be used. To subscribe to an interesting NG, select it from the list click on the subscribe button alongside. After subscription, the NGs show up under the server name. If needed, all the new messages in all the groups can be downloaded simultaneously by selecting the server and clicking on the Get Messages link under the NG menu. Alternatively, messages from the groups can be individually downloaded by selecting the relevant NG and using the Get Messages link.



The Xana News interface

The freshness and number of messages to be downloaded can be set in the Get Messages windows. By default headers and message bodies of all the new messages are downloaded. Once the messages are downloaded, they are displayed, just as e-mails in an e-mail client. The reading and replying, and creating new message functions are similar to that as seen in e-mail clients. Replies can be made to the NG itself or to the e-mail of the particular user. To reply by e-mail the mail account needs to be set. If only headers were downloaded in the Get message operation, the message bodies can be downloaded separately by selecting the relevant messages and clicking on Get message bodies link under the Messages menu. All the menu operations can also be accessed from the context menu. The Quick Get option can be used to directly get messages according to the Get options specified in the settings.

To block any user, to prevent newer messages posted by that user from being downloaded, select the message and click on the Bozo Author button. If there are discussions that are going on which do not interest you they can be marked as Ignored so that newer messages in that thread are not downloaded the next time the "Get Messages" command is given. The Search option can be used to find messages matching certain criteria.

In case many NGs are subscribed to, to update messages in selected NGs use the Batches option in the Tools menu. Here groups of NGs can be operated as a batch, and the common options can be specified. So, the next time these NGs need to be updated, the Batches menu can be used. Alternatively, the Batches toolbar can be activated from the View menu for easier access to the batch controls.

If the discussion in certain NGs have grown stale, the updating of these NGs can be stopped by making them Dormant. In this state, Get Messages commands applied to the server or batch as a whole do not apply to the NG. The NG will have to be individually updated.

To quickly get to an interesting topic it can be bookmarked. The Bookmark pane needs to be activated from the View menu. This shows up below the message list. Dragging and dropping a message to this pane creates a bookmark to that message. Many bookmark categories can be created, and double-clicking on any

entry takes you to the message in that particular NG.

Configuration

Xana allows a rule hierarchy to be used—starting from a client level rule set which can be modified to apply to each NNTP server in the client, which can be further modified to individual NGs in a NNTP server.

A Client level rule can be set from the Options link under the tools menu. Here various options like the appearance of the interface, the connection settings, the shortcut keys, etc., can be set. Under the Default settings link, you will find settings that can be further tweaked at the server and newsgroup level.

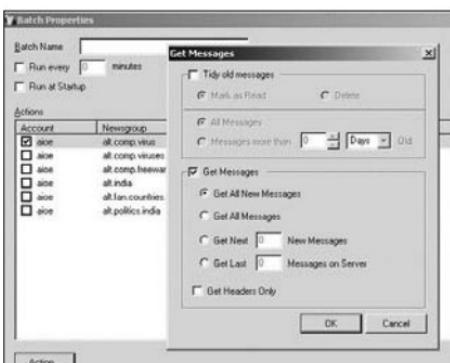
Here, under the User Preferences link, you can specify the default operation when the Get Messages command is given—the number of messages to be downloaded, whether to download message bodies as well or just the headers, etc. Under the Sorting and threading link the default sort order can

be specified. The Posting link can be used to specify the template for your responses, which frees you from the need to type out the additional text. The Filters link allows creation of filter rule to prevent messages from undesirable posters from being downloaded. Since SPAM is also an issue in NGs, this feature is quite useful.

The Mail Accounts link under the Tools menu can be used to specify the outgoing e-mail server to be used when responding by e-mail to any NG user.



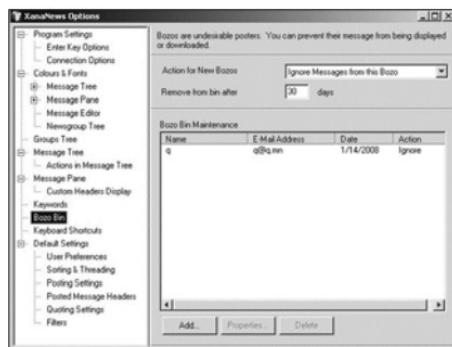
Xana search filters



Batch operations in Xana

Comments

Xana is an impressive NG reader. Displaying the number of topics/messages alongside the list of NGs in the server would aid identifying the level of activity in an NG without having to subscribe to it. But this is a very minor complaint.



Configure your Xana options

7.1.3.2 40tude Dialog (4D)

Usage

At first launch, a wizard will guide you in configuring the application. For some strange reason, the application will not accept domain names with 2 letters only. Besides that minor irritant, the installation is smooth. If needed, 4D can also be used as an e-mail client and the relevant information like POP3 server address, etc., can be entered. The wizard will ask to download all NGs in the server when the configuration is over.

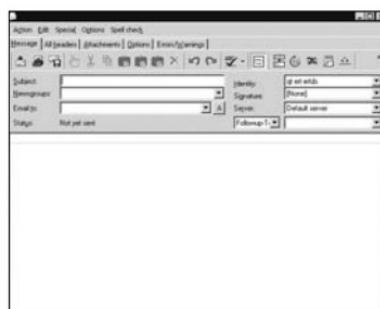


The 4D interface

The list of NGs in the server is displayed along with the new messages in each under the All tab in the top left pane. The message count makes it easier to identify the more active NGs, since there is nothing worse than ending up subscribing to a ghost NG. Under the Filtered tab, you can shortlist NGs based on the presence of the search term in their names. You can sample the headers in an NG without subscribing to it. To subscribe to an NG a double-click suffices. Headers in an NG can be downloaded by selecting the NG and using the Get New headers in the

Selected/Subscribed group link under the Online menu. An alert pops up asking permission to download headers in case there are too many of them. The threshold can be configured under the Default Group options under the Group menu. The downloaded headers are displayed in the right pane. Here you can choose headers to download their message bodies. Double-clicking on the header will download the body, alternatively a group of headers can be selected or marked (from the context menu) and their bodies downloaded simultaneously by using the Get Selected/Marked message bodies option under the Online menu.

To reply, use the Follow to Newsgroup option in the context menu of the selected message. This will cause a message creator window to pop up containing the quoted text. To start a new discussion, use the New Post to newsgroup option in the Post menu. In the same menu the option to send an e-mail is also available.



4D's e-mail client

4D employs a scoring scheme, which helps you quickly identify responses from selected users. A rule can be set to award a score to certain users. The score is shown in a different column along with the header, making it easier to identify responses from that user. Another tool used by 4D to make it easier to avoid certain users is the Plonking scheme. A Plonk rule can be set to delete messages from a user. This way you can avoid spam messages. Unfortunately, the rules cannot be set with a click and drag, and needs to be manually fed in, which can be confusing. Additionally, while the rules can be set from the Context menu, to apply them the relevant option under the Group menu needs to be accessed.

4D also has navigational shortcuts that can be accessed under the Navigate link, which make NG usage easy. NGs can be categorised by using the Category link in the context menu, and the category appears as a tab in the top left pane, alongside other existing tabs like All, Subscribed etc. The usual message handling

options like marking read and ignoring etc can be accomplished from the context menu.

Configuration

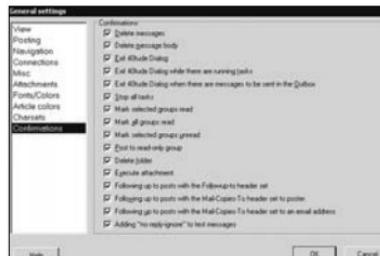
All the configure options are accessible under the Settings menu. The General Settings link offers the options to configure the client. The range of configurable options is daunting, it could easily satisfy someone obsessed with tweaking an interface. The less confusing options relate to setting the text colour and font styles for the messages and interface.

The layout is modifiable from the layout option. A lot of possibilities are presented for you to choose from. Besides changing the shape of the panes, the content in each also can be changed.

The Servers, Identities and Signatures link allows you to set a signature that is appended to all posts. If needed alternative identities can be created, if only to avoid spamming when posting to NGs. The main identity can be used for e-mail purposes. Additional NNTP servers can also be configured here.

Under the Configure Buttons and Shortcuts link the button placement and selection on the toolbars can be modified. It needs to be noted that the interface, though loaded with buttons has them in the wrong place. The context menu in 4D is mostly cosmetic, since the truly useful and much repeated actions can be managed only by accessing the buttons on the toolbar. This makes for a lot of scrolling back and forth, which is an irritant.

The Glossary can be used to create word associations so that typing a code will result in the much larger phrase being pasted automatically. The level of automation can also be taken up a notch by the inbuilt scripting support, accessible from the Scripting link. Frequent actions can be scripted, and the scripts can be launched by certain events. The comprehensive help file that is included has a list of all commands and events that can be used for scripting.



4D's nit-picking options

The Default and Selected Group Options link allows setting the parameters to handle groups. Here various settings like the purge interval, the identities to be used, the message sorting order etc can be used.

Comments

4D is a commendable effort. Its list of skills is long. Support for different encoding formats, multiple server support, multipart message support, default message handling etc are impressive. The only grouse that any user will have is the weak right-click, which has been grossly underutilised.

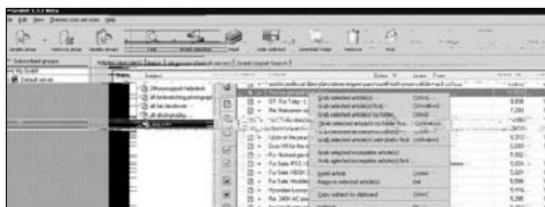
7.1.3.3 GrabIt

Introduction

This is not exactly a NG “reader”. It’s sole function is to download attachments (or binaries) that accompany messages posted to NG. And that it does well, with minimal fuss.

Usage

After installation, at first run you has to enter details of the NNTP server and login details where necessary.



GrabIt's focussed interface

GrabIt will ask to download the NGs in the server. After the NGs are downloaded, they are listed in the right pane of the interface.

The message count is also thoughtfully listed. Right clicking on the NG will offer a Hobson’s choice—“Subscribe”. A subscribed NG is listed under the server name in the left pane. To get the headers, a right click on the NG name will offer the options. A Full Update will download all the headers from the server, an Incremental update will only download the new ones. Unfortunately there are no user limits that can be applied to the number of headers that should be downloaded. After the headers

have been downloaded (which could take time depending on the number of messages), they are listed on the right pane. The message details displayed include the topic subject, the sender, and the number of lines, but most importantly the size. Large message sizes are indicative of the presence of downloadable attachments (since there are no other indications of the presence of an attachment). Right-clicking the message will offer the option to "Grab selected article". Selecting this will download the message—if it is a text message a text file is saved, else the attached file is saved in the download folder. Grabbed messages are deleted from the queue, lest a user download the same message again. A user can read the message by clicking on the Read Article link in the context menu. But there are no options to reply.

The Batch tab in the right pane shows the messages that are being presently downloaded. A NNTP server can be searched for interesting groups by using the "All Groups" tab. GrabIt also allows searching for NGs online, by integrating the Usenet search tab (though, this requires you to register at shemes.com, for a price). Messages in an NG can be searched by clicking the Find button in the toolbar menu.

Configuration

The configuration options can be accessed from the Settings link under the Edit menu.

Here the default download folder can be specified under the Folders tab. Under the Repair and Extract tab GrabIt's archive repairing ability can be enabled. The multiple parts of a file have a par extension. The downloading and error checking of the files can be fine tuned under the Advanced tab.

Comments

GrabIt is an excellent tool to grab binary files off Usenet. It's support for Multi part files and inclusion of tools to repair corrupt



GrabIt's host of options

RAR files make it attractive for those doing a lot of downloading. The lack of even basic facilities to respond to text messages makes it absolutely worthless for those who do no downloading.

7.1.4 Other Clients

We shall briefly discuss the NG handling capabilities of two popular e-mail clients—Outlook Express and Thunderbird.

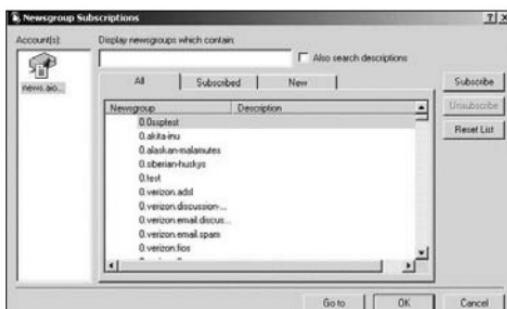
Outlook Express 6

In Outlook Express (OE), a News Account can be added just like adding an e-mail account. Relevant details about the server should be provided. After the setup, you are asked whether to check for NGs. The list of NGs is then provided. A user can search for NGs having specific terms from the Search box. The lack of a message count is disappointing. Right-clicking on the selected NG allows you to subscribe to it. Subscribed NGs appear in the left panel under the NNTP server name. Selecting the NG will cause the message headers in it to be downloaded. Selecting a message header will cause the message body to be downloaded. Neat.

Message management options are similar to those seen in other clients—messages can be marked read, and clicking on reply will pop up the familiar message creation window. In case of multi part messages, selecting the parts and right clicking will reveal the “Merge and Decode” option which recreate the original file from the parts.

The maximum number of headers to be simultaneously downloaded can be set from the Options under the Tools menu. The automatic downloading of the body when selecting the header can be disabled here.

OE is capable of handling multi part messages and allows filters to be set to check the inflow of spam. The folder where messages are stored can be compressed to free space



Creating a news account in OE is easy

by accessing the NG's properties in the context menu.

Overall, OE is more than adequate for most NG users. Familiarity is one of its strongest assets.



You see your news in the good old OE interface

Thunderbird 2

A news account can be created in Thunderbird (TB) from the New Account link in the File menu. After server and user details are provided, the familiar TB interface is shown. The server is shown on the left panel. To get the NGs in it select it and use Subscribe link in the context menu. After the NGs are listed, you can search for interested NGs from the search box. There number of messages in the NG is not listed which is a disappointment. To subscribe to the NG, select it and click on the subscribe button. The NGs are listed under the server name in the left panel. Selecting the NG will cause the headers to be downloaded. A pop up will prompt you if the number of headers to be downloaded is huge. After the headers are listed, selecting them will automatically download the message body. Replying to and creating posts will launch the usual message creation window.

Message filters can be created and applied, just as in the case of e-mails.

To configure server level options, right click on the server name on the left panel and use the Properties link. The client behaviour can be modified from the Options link in the

Tools menu. Finer control can be exercised by accessing the Config Editor under the Advanced button. Tip: to quickly get to controls that pertain to NGs use the search box.

Overall, TB offers all the required functionality to satisfy most NG users.

7.1.5 Conclusion

NGs offer a near-academic environment to indulge in serious discussion. When one tires of all the eye candy on a Web forum, and

is ready to focus on the brasstacks of serious exchange of ideas, an NG is the place to be. A news-reader that can make the process of finding compatible NGs easy, and handle the



The Thunderbird interface

basic functions of reading and replying to messages is a valuable tool. While standalone apps exist, whether your needs are met by an extant e-mail client or even by sites that allow NG access needs to be evaluated before proceeding. The only real need to have a standalone app seems to be when dealing with binaries.

7.2 RSS Feed Readers

7.2.1 Introduction

Syndication is, simply put, the practice of offering content to publishers. This is commonly seen as cartoon strips in newspapers, where the artist syndicates his work to newspaper publishers. More than one newspaper can carry the same syndicated cartoon on the same day. The online version of syndication works on similar lines: someone who creates content—like news sites or bloggers—allows the content to be passed on to those who would like to publish it on their site, or just be informed of new posts.

The message that is sent out is called a feed, which besides the core content of the creator also contains meta information. There are mainly two formats in which the feed is structured—RSS and Atom, both of which are based on the XML file format. RSS or Really Simple Syndication (or Rich Site Summary) is the more commonly used format (there are different versions of RSS including one called RDF) and it is common to refer to a feed as an RSS feed even though it may be an Atom feed.

To be able to display the feed, one needs an application that can recognise the main formats. A standalone application is called

a feed reader or feed aggregator. Newer browsers like Opera 9 and IE7 include an RSS feed reader, as do some e-mail clients. Additionally there are Web sites that can offer similar functionality as a standalone client.

Websites that create RSS feeds have links to them on the site. If the browser supports RSS feeds, clicking on the link will launch the inbuilt client or the external feed reader.

The list of sites offering a feed can be had from many online sources, like syndic8.com. Some sites offer a list of feed links as a file. Such a list is created in OPML (Outline Processor Markup language) which is also based on XML, and all feed readers are capable of reading/opening it. Lists of feeds created by you can also be saved as an OPML file which makes it easier to open these feeds in another feed reader.

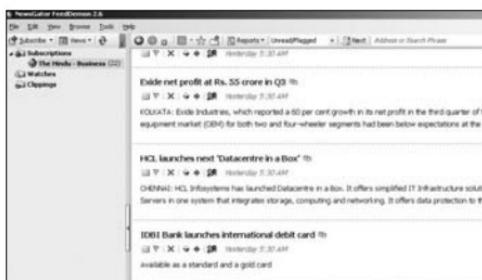
7.2.2 Feed Readers

We shall discuss a few feed readers next.

7.2.2.1 FeedDemon

This comes across as a thoughtfully, well put together feed reader. After installation, you have the option to load the preinstalled feed lists. If needed, a user can subscribe to a new feed by using the Subscribed button. After entering the XML address, you have the option to preview a feed before continuing with the subscription. During the installation, you are presented with a form to sign up at Newsgator.com for free. As a member of the site, you can also search for feeds on any site by just entering the Web site address. In practice, though, this didn't seem functional.

The layout of the panes can be changed under the View menu. FD allows messages to be categorised by flagging them or moving them to the Clippings



The FeedDemon interface

folder which is accessible under the list of subscribed feeds. The right-click also offers option to send the news item directly to a blog, which first needs to be configured with data about the server. Presently, this is an uphill task. The news item can be sent to bookmarking sites like digg.com and del.icio.us.

Prefetching is the process of downloading news items so that they can be read offline. This can be done from the file menu. You can also search for news items from the search bar on the top of the pane.

The Feed Properties link can be used to configure the handling of the NG by the client. The number of headers to download and the interval for checking for new news items can be specified here. The Options link under the Tools menu allows configuring the client behaviour. Here options like Automatic updation of feeds can be disabled. Other options that can be configured are the Appearance aspects, the browser to be used to view a news item in full, the kind and number of alerts to be generated, etc.

The Clean up wizard in the tools menu can be used to delete messages meeting the criteria mentioned in the wizard. The Hit the Panic button can be used in case the number of messages is too high. Clicking on the button changes the status of messages matching the set criteria to read. Those who have registered for a newsgator account can compare the contents of the folders offline and online. These folders can be synchronised by the Synchronisation options link under the tools menu.

7.2.2.2 Abilon News Aggregator

Abilon News Aggregator (ANA) is a fine looking and capable application, that is simple to use.

After installation, you are presented with a preconfigured list of feeds that can be readily subscribed to without hunting for online sources. It is also capable of loading feeds from local XML files or OPML files or remote URLs too. It can also export the list of feeds to an OPML file. The layout can be modified from the default three vertical pane layout. Feeds can be viewed in the right-most pane. Clicking on the subject takes you to the Web site within the app itself. It implements a tabbed interface with each link opened

in a different tab. News items can be marked as persistent so that they are retained even when the rest of the news are deleted. An alternative is to classify interesting news items under memories. These news items are made visible under the Memories pane on the left, and also avoid deletion when the feed is cleared.



The ANA default interface

ANA also includes a blog client that can allow direct posting into blog hosts like blogger.com and livejournal. You need to provide the server details (which we could NOT find easily) before the post is published in the blog. Once done, though, you can blog a news item with just a right-click.

One area where ANA fell short was the inability to integrate smoothly with Firefox—something Feedreader (up next) did easily.

7.2.2.3 FeedReader

FeedReader (FR) is unwieldy to work with, though it is loaded with features. The only point in its favour is its ability to integrate with Firefox. When using Firefox, on clicking on the XML button, the option to use any of the installed feed readers or Firefox's inbuilt feedreader is presented. FR manages to make its presence felt in that list. And clicking on the subscribe button launches FR with the URL listed.

FR offers all the basic controls to manage feeds—namely opening local or remote XML or OPML files. It can also export the list of feeds it is



FR's default interface

working with in OPML format. Clicking on any link in the feed's news item opens the page in the application itself.

It carries a lot of navigational aids that make managing news items easier. News items can be grouped on their source, date or on their tag. A tag can be applied by you by right-clicking on the news item. Messages can be made hidden based on their read status. It regularly checks for newer items in the feed, and this information is displayed near the system tray as a pop up.

A Smartfeed is a feed that can be created from the other feeds that are being handled by FR. Criteria for including news items from other feeds into the smart feed can be specified. So accessing the smartfeed will list those feeds that match the criteria. To create a new Smartfeed use the Net button. Smartfeeds use a blue icon in contrast to the orange icon of normal feeds. The Smartfeed appears alongside other feeds in the left panel, and can be named.

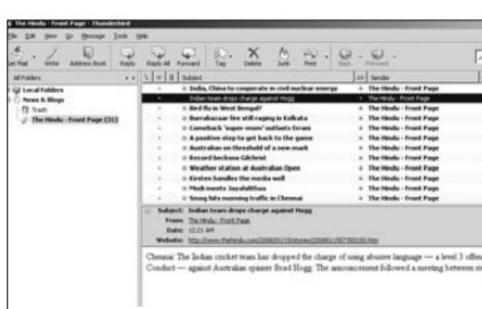
Under the Tools menu, the Preload Feeds link opens a list of preloaded feeds which are categorised into different groups, from which you can easily select and subscribe. Feeds that have additional content attached can be handled with the Enclosure browser which allows downloading the attached content. The Options link allows configuring FR.

Overall, FR comes across as a capable feed reader, but it somehow lacks intuitiveness, which can be discouraging.

7.2.3 Other Feed readers

Thunderbird

This jack of all trades also includes a feed reader. After the local or remote feed address is provided, the familiar three pane window is displayed. The news headings are displayed on the top, and select-



Thunderbird's feed reader interface

ing them will reveal the entire item in the bottom part. To visit the Web site, the link is presented in the header. TB allows tagging the news items—which is a feature available for other services like e-mail as well.

Thunderbird is a very basic feed reader that can be of use to those who already use it for e-mail or newsgroups.

7.2.4 Online Feed Readers

There are many online sites that allow a user to view feeds from multiple sources. In many cases getting registered on the site allows the feeds to be customised. A long list of online Feed readers is available at <http://www.newsonfeeds.com/faq/aggregators> and is worth a look. Grabline.com, which is mentioned in that list, is of particular interest for Indians, since it brings together feeds from the major news papers like Indian Express, Times of India, Hindu, Telegraph, etc., besides a few other Indian Web sites. The number of irrelevant stuff on the page could be distracting though.

Google Reader

This is the feed reader offered by Google. By extensively using AJAX the interface is made extremely responsive and intuitive. You need to create an account at google.com to avail of the services. Once signed in, accessing google.com/reader launches the service. Google recommends a few feeds which you can ignore. The Add subscription link allows you to search for feeds matching terms. To subscribe to any of the feeds in the search results just click on the Subscribe link. The news items in the feed are displayed by selecting the feed name in the left panel. The interface allows changing the display properties, and sort order. Clicking on any of the news item links causes the Web site to be loaded in a different browser window/tab. The interface allows other options like e-mailing the news item or link, tagging the news item for easy retrieval, etc.

The experience of using the reader service is a pleasant one. The biggest drawback of online feed readers , which also includes Google's service, is the inability to view the news items offline.

7.2.5 Conclusion

For a prolific Web user, feeds can be your saviour. Rather than visit the many sites to find updates, having them send their feeds to you is more convenient. With all types of feed readers readily available, it is up to you to choose which features are most important, and settle for a client or online service that offers it.